VERIEQUIVBENCH: AN EQUIVALENCE SCORE FOR GROUND-TRUTH-FREE EVALUATION OF FORMALLY VERIFIABLE CODE

Anonymous authorsPaper under double-blind review

ABSTRACT

Formal verification is the next frontier for ensuring the correctness of code generated by Large Language Models (LLMs). While methods that co-generate code and formal specifications in formal languages, like Dafny, can, in principle, prove alignment with user intent, progress is bottlenecked by specification quality evaluation. Current benchmarks rely on matching against ground-truth specifications, a manual and expertise-intensive process that has limited existing datasets to a few hundred simple problems and also suffers from a reliability issue. To address this, we introduce VeriEquivBench, a new benchmark with 2,389 complex algorithmic problems that probe the limitations of current models in both code generation and formal reasoning. Our evaluation framework replaces ground-truth matching with a formally grounded metric, the equivalence score, and rigorously verifies the quality of generated specifications and code. Our results show that generating formally verifiable code remains a profound challenge for state-of-the-art LLMs. This underscores both the difficulty of the task and the need for benchmarks like VeriEquivBench to drive progress toward scalable and reliable coding agents.

1 Introduction

Large language models (LLMs) already possess substantial capacity for following natural-language instructions and executing a wide range of coding tasks (Li et al., 2022; Jain et al., 2024; Zhao et al., 2025). At the same time, the correctness of the generated code remains a concern (Cotroneo et al., 2024; Wang et al., 2025b), where functional errors cost users extra effort to debug and also pose significant risks in the safety-critical domain (Dalrymple et al., 2024). A common solution is to evaluate generated code through unit tests (Jimenez et al., 2024; Wang et al., 2025c). However, this process offers no provable guarantee of correctness, as insufficient unit test coverage can fail to detect critical errors (Yu et al., 2025). On the contrary, a verifiable system resolves the issue by co-generating formal specifications and code to formally verify the alignment with the natural language query intention (Sun et al., 2024). Our work focuses on building an end-to-end agent for formal verification, for which we adopt Dafny (Leino, 2010). It is an ideal choice as Dafny's automatic theorem prover (De Moura & Bjørner, 2008) eliminates the need for manual proof writing. Furthermore, its similarity to common languages like Python and C simplifies code transformation.

While several benchmarks (Ye et al., 2025b; Thakur et al., 2025) target at building a reliable reasoning system by formally ensuring the exact code generation (György et al., 2025), their progress is constrained by the reliance on manually-written, ground-truth specifications for evaluation. This formal annotation process is incredibly labour-intensive and requires deep expertise (Misu et al., 2024), which sets a barrier to scaling these benchmarks in both size and complexity. As a result, prominent Dafny benchmarks, including DafnySynthesis (Misu et al., 2024) and CloverBench (Sun et al., 2024), contain only 215 simple examples combined, insufficient for evaluating current LLMs' advanced reasoning abilities. Moreover, the reliance on expert annotation is not only a scaling bottleneck; it also leads to a reliability issue. An analysis (Sun et al., 2024) has figured that 10% of expert-written specifications in DafnySynthesis are wrongly claimed as ground-truths, and our own review finds another 18%, containing errors or ambiguities. Such flaws undermine the validity of any benchmark that depends on a ground-truth solution. This raises a critical question: *How can*

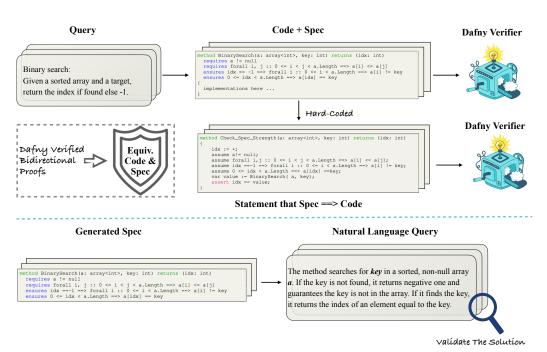


Figure 1: An end-to-end verifiable coding agent first generates code and specifications, using the Dafny verifier to prove their mutual equivalence. It then translates the complete formal specification back into natural language, allowing the user to confirm that it aligns with their original intent.

we reliably evaluate specifications' quality without depending on the ground-truth? To answer this question, we make the following concrete technical contributions:

Contribution 1. We propose a novel formally-grounded metric, named the *equivalence score*, that measures the mutual equivalence between generated code and its specifications. The score confirms whether a specification unambiguously describes the code's behaviour by using the Dafny verifier to check for bidirectional implication. This automated process has no false positives, ensuring that only correctly matched code-specification pairs are accepted. In order to validate the alignment with the query intention, we further include a second evaluation step: translating formal specifications back to natural language, as illustrated in Figure 1. Note that the effort in translation between natural and formal languages is asymmetric, since natural language can be ambiguous and lack necessary logic (Jiang et al.). Using Claude-4 as a judge (Wang et al., 2025a), we observe a high success rate of 82.98% for Grok-4 translations, confirming its viability as an evaluation metric.

Contribution 2. Equipped with our automated evaluation metric, we introduce *VeriEquivBench*, a benchmark of 2, 389 examples for end-to-end formally verifiable code generation. VeriEquivBench significantly expands on prior work in both dataset size and problem complexity, a leap demonstrated by the average Cyclomatic Complexity score, which rises from 2.44 in DafnySynthesis to 5.63. The core of our dataset is converted from the LeetCode corpus (Xia et al., 2025), a large and community-validated collection of algorithmic problems well-suited for evaluating a model's reasoning abilities. To supplement this data, we also introduce a synthesis pipeline that uses a structured tagging system to generate novel queries by randomly combining tags for different domains, data structures, and algorithms, introduced in Section 2.3. This provides a scalable method for creating large training datasets of new problem descriptions that are fully compatible with our automated evaluation signal. However, Claude-4 is able to generate qualified Python code for only 15.85% of these novel synthesis queries, reflecting the model's limitations.

Contribution 3. We conduct a concrete evaluation of state-of-the-art LLMs, where VeriEquivBench serves as a testbed for these models to explore and extend the reasoning abilities on complex problems, beyond human-annotated data (Silver et al., 2021; Ye et al., 2025a; Shojaee et al., 2025). Our evaluation highlights the profound difficulty of this task and the effectiveness of our benchmark. The best-performing model, Claude-4-sonnet, which solves 75.81% of the problems in CloverBench, succeeds on only 4.83% of our data, even with a pass@4 metric. Given this poor performance, and following prior work (Loughridge et al., 2025; Yan et al., 2025), we include two simpler auxiliary

tasks to scaffold the problem: (1) infilling the necessary clauses to complete a given specification, and (2) generating a specification from provided Dafny code. To facilitate future research, we establish baselines for both tasks using reinforcement learning.

2 BENCHMARK OVERVIEW AND CONSTRUCTION PIPELINE

In this section, we first present aggregate data statistics for VeriEquivBench. Subsequently, we introduce the two curated subsets released with the benchmark: (i) the LeetCode-transformed dataset, and (ii) a tag-composition dataset, called TagComp, the latter being explicitly constructed to evaluate verifiable agents on novel data without contamination (Tu et al., 2024; Riddell et al., 2024).

Each problem in our benchmark provides a comprehensive set of artifacts: a natural language query, implementations in both Python and Dafny, unit tests and two versions of formal specifications: a strong auto-formalized baseline explained in Section 2.1 and a weaker, verifiable but incomplete version explained in Section 2.2. Additionally, each problem is annotated with metadata, including its difficulty level and descriptive tags for the relevant algorithm, data structure, and domain. Unlike LeetCode, our benchmark uses a more detailed and structured set of tags to categorize problems. This new tagging system is described in Section 2.3 for future query synthesis.

Starting from the original Leetcode split of 2,641 training and 228 test instances, we first curate 2,174 cases successfully transformed to Dafny. Then we compose new problems by merging tags, producing 1,893 additional items; the full tag-composition procedure is described in Section 2.3. For new problems, we ask Claude-4-sonnet to generate pairs of Python code and corresponding unit tests. For only 300 of new problems, Claude-generated code passes at least 85% of their corresponding unit tests, forming the cleaned TagComp dataset. Of these, 215 samples clear the weak-baseline pipeline, giving us 2,389 problems in total that pair natural-language queries with formally annotated code.

Table 1 presents key metrics for our annotated Dafny code, which uses the weaker, verifiably correct specifications. Our problems are significantly more complex than those in CloverBench, often involving multiple methods rather than a single one. Furthermore, the corresponding specifications, while incomplete, contain a substantial number of formal clauses.

DatasetMetricfunctionmethodinvariantensuresdecreasesLeetCodemean0.781.335.121.710.46

3.18

7.34

3.14

0.70

0.96

Table 1: The table overviews several attributes of our annotated code.

2.1 LEETCODE AUTOFORMALIZATION

mean

TagComp

Past formal-language sets such as DafnyBench (Loughridge et al., 2025) are still small and narrow, because hand-written specifications are too costly to scale (Misu et al., 2024). To obtain large, varied training data without extra human cost, we mine the classic Leetcode pool, convert problems to formal specifications, stated in Figure 2 Pipeline 1, while keeping query and specification aligned with two short tightening evaluation protocols (Sun et al., 2024), shown in Figure 2 Pipeline 2.

Specification Generation We feed the problem description to Claude-4-sonnet to obtain an initial Dafny specification, yet even the initial drafts often contain syntax errors. Thus, we revise and resubmit up to ten times until the file has no parse or resolution errors. We find that supplying two simple examples exploits the model's in-context learning (Dong et al., 2023) and sharply lowers the error rate (prompt template in Appendix C).

Furthermore, we constrain the model to generate specifications using only first-order logic, prohibiting recursive or dynamic programming-style definitions. This ensures the specification describes the problem's declarative properties without leaking the implementation's structure.

Equivalence to NL The equivalence check follows the protocol proposed by Clover (2024) and contains two steps: (1) A model (we use Grok4 here) rewrites the description so that it cleanly

mirrors the specification, then another model (Claude-4) judges the equivalence between the original description and the rewritten one, yielding a score; (2) The specification alone is translated into Python and executed against the ground-truth LeetCode unit tests. The unit test passing rate is reported in Appendix D.1.

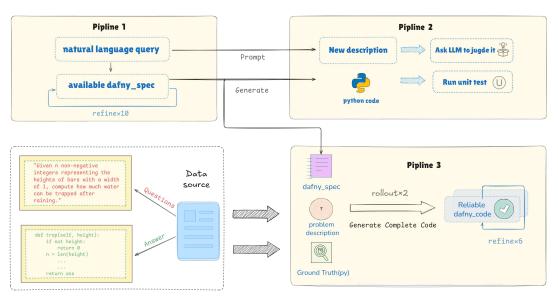


Figure 2: The figure outlines our autoformalization and code generation workflow: Pipeline 1 produces comprehensive and syntax-free specifications; Pipeline 2 checks consistency between the NL query and the specifications; Pipeline 3 emits fully annotated code that passes the verifier.

2.2 LEETCODE VERIFIABLE CODE GENERATION

Owing to the limited performance of state-of-the-art LLMs on challenging formal-language coding, we adopt the multi-stage pipeline (pipeline 3 of Figure 2): prompted by the previously generated specifications, the problem description, and a reference Python solution, the stronger model (Claude-4) produces annotated Dafny code, while a lighter model (Claude-3.5) then polishes this output, iterating up to six times to eliminate any syntax or parsing errors.

In practice, the vast majority of problems converge within three refinement rounds, while a residual subset still fails to yield a well-formed artifact even after the sixth attempt, with the success transformation rate reported in Table 2, and data statistics reported in Table 1.

Table 2: The table shows the number of examples without syntax errors in autoformalization and verifiable examples in code generation.

Dataset	Spec Autoformalization		Dafny Code Genetation	
2	Number	Rate (%) Number		Rate (%)
LeetCode	2584	90.1	2174	75.8
TagComp	296	98.7	215	71.7

2.3 Data Synthesis Through Tag Composition

We propose constructing a fine-grained, "template-level" taxonomy to provide an abstract description of algorithmic problems via tags (Wang et al., 2025c). In our system, every task is labelled with three orthogonal categories: **domain**, **data-structure**, and **algorithm** class((Chollet et al., 2025)).

To obtain these labels, we (i) harvest a high-quality seed pool from the Luogu online judge (luo, 2025), and (ii) manually prune hallucinated or off-topic tags. Our ontology defines over 500 fine-grained tags, offering more than seven times the descriptive granularity of the 69 tags used by Leet-

Code (see Appendix A for a comparison). The tag set is designed so that, taken together, the tags collectively reflect the complete programming knowledge entailed by each individual problem, while retaining a modest level of abstraction.

The three categories of tags capture complementary aspects of programming knowledge. First, the domain category encompasses the overarching problem space or application context in which an algorithm operates, such as graph theory. Second, the data structure category pertains to the foundational mechanisms for manipulating data that underpin the algorithm's functionality and efficiency, like arrays. Third, the algorithm category refers to the core strategic paradigm employed, such as sorting, focusing on the decision-making logic. These algorithm tags directly shape the overall control flow of a solution, as they orchestrate the program logic and structure.

However, not all problems conform to highly standardized patterns. In contemporary algorithmic competition problems, for instance, many challenges necessitate solvers to discern the underlying mathematical structures, an approach commonly termed "constructive methods". From a coding perspective, these constructive methods typically appear as compact code blocks that rely solely on fundamental loops or arithmetic operations. Consequently, it is difficult to categorize them beyond a general "constructive method" tag. From the problem setter's viewpoint, such problems and their solutions stem from empirically observed mathematical structural properties, which inherently resist exhaustive coverage by conventional tags.

To synthesize novel queries, we select tags in the following workflow: first, we randomly pull 12 tags from each of three pools, and then let Claude-4 pick any 3–8 tags in total. This short list is fed back to the model so that Claude can create one clear algorithm question with roughly 40 unit tests (Xu et al., 2025). Initially, we create approximately 1,900 problems, but only retain the 300 that pass at least 85% of their tests (Xu et al., 2025), and call this clean set, *TagComp*. The detailed pipeline and prompt templates used can be found in Appendix B and Appendix C.

3 EVALUATION METRICS AND TASKS

A verifiable coding agent reduces hallucinations and provides trustworthy solutions aligned with users' intentions. As shown in Figure 1, our solution evaluation includes two steps, which are

- verifying the equivalence between generated code and specifications, and
- validating the solution by translating formal specifications back to problem descriptions in natural languages.

To understand the need for verifying the **equivalence** between the code and the specifications, consider a simple binary search algorithm. The goal is to return the index of a key in a sorted array a, or negative one if the key is not found. A weak but verifiable post-condition might only state that the output, idx, is within a valid range: ensures $-1 \le idx \le a$. Length. While this specification passes the verifier, it fails to exactly describe the code. This creates a dangerous loophole: an incorrect implementation that doesn't actually find the key could still satisfy this weak condition, and the verifier would not catch the error.

Existing benchmarks do not offer a metric to formally validate the quality of specifications. Without one, there is no way to guarantee that the verified code truly aligns with its intended behaviour. Instead, building up equivalence examines whether the specification is complete without ambiguities. Our equivalence score accomplishes the task by proving the bidirectional implication relationship:

- whether the code falls into the lattices described by the specifications, and
- whether the method output is the unique value satisfying specifications for any inputs.

Both proofs can be automatically completed by the Dafny verifier. The first direction can be verified by passing the annotation to the verifier. The second direction requires creating a statement that the specification implies the code for the verifier to check.

Figure 3 presents a counterexample to illustrate how our equivalence score identifies an underspecified function. The Max method correctly returns the maximum of two integers a and b, but its post-condition (ensures max $\geq a$) is too weak; it doesn't guarantee that the output is also greater than or equal to b. To test if the specification fully implies the code's behaviour, we use the Check_Max_Spec method. This method creates an arbitrary value max, assumed to satisfy

```
method Check_Max_Spec(a: int, b: int) returns (max: int)
{
    max := *;
    assume max >= a;
    var value := Max(a, b);
    assert value == max;
}
```

Figure 3: We show an example where the equivalence score proves the given specifications are underspecified for returning the maximum between two integers. The code presents the statement to verify whether the specification implies the code.

all provided pre-conditions and post-conditions. Our equivalence score then tests the assertion that the unique value satisfying these conditions is the method output. The Dafny verifier is guaranteed to find this assertion to be false without any false positives. Because the specification is not strong enough to imply the code, this program would not receive an equivalence score.

As mentioned in the introduction, end-to-end formally verifiable code generation is still challenging for current proprietary LLMs. Dafny has its own programming logic, such as claiming the invariance of old elements in arrays to support the proof. Therefore, we re-emphasize the importance of two auxiliary tasks to facilitate understanding of specific nuances of Dafny, introduced in DafnyBench (Loughridge et al., 2025) and Veri-Code Series I (2025):

- **Verifiable Code Refinement:** Given annotated but unverified Dafny code, the model's goal is to add the necessary intermediate clauses, such as invariants and lemmas, to make the code pass the verifier. Success is determined by successful verification.
- Code-To-Spec Generation: Given a Dafny implementation, the model attempts to generate the strongest formal specification. The quality of the output is evaluated by measuring its strength improvement over a baseline, using the spec-superior-score (Yan et al., 2025).

Our two sets of formal specifications map onto these auxiliary tasks. For the Verifiable Code Refinement task, models are challenged to fix our strong auto-formalized specifications. For the Code-to-Specification Generation task, models improve upon our weaker, but already verified, specifications.

4 EMPIRICAL EVALUATION

This section validates the quality of our benchmark and the reliability of our evaluation metric. We then present the performance of several state-of-the-art LLMs on the end-to-end verifiable code generation task, followed by an analysis of our baselines on the two auxiliary tasks.

4.1 QUALITY METRICS

Specification Quality Our strong specification baseline, generated via auto-formalization, contains the ground-truth specification for 7.14% of the LeetCode-derived problems and 7.87% of the synthetic TagComp problems, shown in Figure 5. In total, this process yields 161 complex algorithmic data with rigorously verified specifications. This significantly enriches the publicly available dataset of ground-truth specifications.

Code Transformation Quality To evaluate the quality of our Python-to-Dafny code transformation, we randomly sample 20 programs from the LeetCode set and validate them against translated unit tests. Due to the mismatch between Python and Dafny unit test formats and the long compilation time of Dafny, we only provide limited validation. But the transformation is proved highly reliable, with 90% of the translated Dafny programs passing all tests.

Data Complexity The average Cyclomatic Complexity (McCabe, 1976) quantitatively manifests the increasing complexity of our data, which counts the number of linearly independent paths in the control flow graph. It is computed using the Radon software package for Python.

We list the score for MBPP (Austin et al., 2021), since 50 manually annotated data in DafnySynthesis are based on MBPP-50 and the other 103 are also transformed from it. Thus, the analysis represents a comparison to DafnySynthesis. We skip the analysis of CloverBench due to a lack of Python implementations. Our benchmark's average score of 5.63 is significantly higher than the 2.44 for DafnySynthesis, indicating more complicated control flows. Notably, our synthetically generated data is even slightly more complex than the LeetCode-derived portion, with a score that is 0.25 points higher. This complexity is further validated by a manual rating from Claude-4, which classified the majority of our synthetic problems as either medium or hard.

Table 3: The table compares the code complexity of a previous benchmark and VeriEquivBench, indicating a more intricate control flow of our data.

Dataset	MBPP-50	MBPP	LeetCode	TagComp
Average Cyclomatic Complexity	2.44	2.78	5.38	5.63

4.2 VALIDATION OF THE EVALUATION METRICS

We first validate our equivalence score on 50 expert-written verifiable code provided in DafnySynthesis. CloverBench has reviewed their data and reported that 10% of the data does not give the ground-truth specification. After testing on our evaluation metric, the equivalence score, we figure out another nine examples where the formal specification contains ambiguities or the original code has errors. An example is shown in Figure 4, where the formal specification does not specify the invariance of array length and leaves a logic gap. However, only eight examples out of 14 failures are successfully fixed by us, demonstrating the hardness in manual annotation. All examples with wrongly claimed ground-truth are listed in Appendix D.4 with the issues stated.

Figure 4: An example of a weak specification in sample #625 that fails equivalence scoring. The formal specification is ambiguous as it omits a post-condition on the invariance of the array's length.

Next, we evaluate all previous benchmarks and observe a serious quality issue in previously provided ground-truth formal specifications, shown in Table 4. It has been discussed that the equivalence check relying on natural language provided in Clover has limitations, and it turns out that a large number of specifications do not establish the equivalence with the code. Meanwhile, DafnyBench is not designed for checking the completeness of specifications and thus, gives the lowest score.

Furthermore, we evaluate Grok-4's translation ability, using Claude-4-sonnet as a judge (Wang et al., 2025a). We test it on our filtered auto-formalized specifications derived from LeetCode and observe a high success rate of 82.98%, validating it as a reasonable evaluation metric.

Table 4: We present the percentage of data gaining the equivalence score in previous benchmarks.

Dataset	DafnySynthesis	CloverBench	DafnyBench
Equivalence Score	76.22%	61.29%	43.09%

4.3 VERIFIABLE CODE GENERATION

Figure 5 (b) and (c) present the pass@4 results of three proprietary LLMs on end-to-end formally verifiable code generation, tested on CloverBench and our contamination-free synthetic set, Tag-Comp. On the previous CloverBench benchmark, a capable model like Claude achieves a 75.81% success rate, with most errors stemming from issues in specification writing rather than fundamental code generation flaws. However, on our more challenging TagComp dataset, this performance collapses. While the model achieves code-specification equivalence on 10.34% of the problems, more than half of these successes are undermined by incorrect code generation, demonstrating that our benchmark effectively tests both coding and formal reasoning abilities.

A closer look at our benchmark results reveals the challenge of verifiable code generation. While Claude is most capable of producing syntactically correct Dafny code, all three models struggle significantly with generating mutually equivalent code and specifications aligned with the query intention. In our rigorous two-step evaluation, the equivalence score measures the formal alignment of code and specifications, while the exact matching score further validates against the original natural language intent. Ultimately, fewer than 5% of the generated solutions from any model pass this framework. This result underscores the difficulty of formally verifiable code generation on complex algorithmic problems, requiring strong coding and formal reasoning abilities.

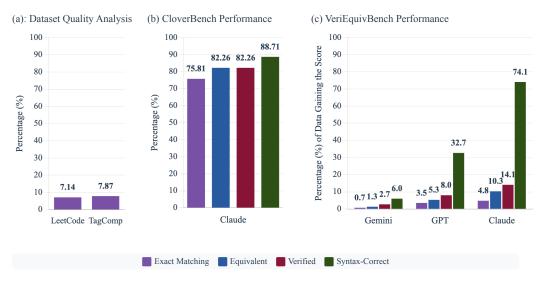


Figure 5: Exact matching score measures the percentage of data passing our two-step evaluation framework, giving solutions aligned with the query intention. Part (a) gives the amount of verified ground-truth solutions in our benchmark. Part (b) shows that the previous CloverBench benchmark is too simple to properly evaluate the advanced reasoning abilities of capable models, as evidenced by a high success rate. Part (c) presents the pass@4 performance of gemini-2.5-flash, gpt-5, and claude-4-sonnet on our end-to-end verifiable code generation task.

4.4 AUXILIARY TASKS

For the two auxiliary tasks mentioned in Section 3, we provide two RL-trained baselines, with the SFT model provided in Veri-Code Series I (2025). As stated, the verifiable code refinement task uses passing the verifier or not as the reward to infill intermediate clauses, while the spec generation task uses the spec superior score introduced by Yan et al. (2025). Spec superior score measures whether

the generation specifications described the code better than our weak baseline. The choice of the RL algorithm and hyperparameters follows their implementation as well. We split our LeetCode transformed data into three parts with 1770 training data, 200 validation data and 204 out-of-domain test data, using tags uncovered by the training data.

Our baseline scores 17.68% for the refinement task and 54% for the spec generation task on the validation set. However, in the spec generation task, almost no data generates a complete specification, resulting in an equivalence score. A possible reason is that the SFT model provided is trained on overly simple problems and does not have enough exploration ability. The training curve and results on the test set are presented in Appendix D.3.

5 RELATED WORKS

A central challenge in advancing Large Language Models is developing metrics that not only assess performance but also provide a clear signal for improvement we desire. **Outcome-based metrics**, such as final-answer accuracy in mathematical reasoning (Cobbe et al., 2021) or pass rates on unit tests in code generation (Austin et al., 2021), are prevalent but limited. They disregard the fidelity of the reasoning process and remain susceptible to false positives, a limitation shared by methods employing external solvers for verification (Huang et al., 2025; Feng et al., 2025).

Formal verification offers a more rigorous evaluation alternative, using proof checkers like Dafny (Leino, 2010) or Lean (De Moura et al., 2015) to provide an unambiguous correctness signal without requiring a ground-truth solution. However, in verifiable code generation, this signal is fundamentally unidirectional: it validates that the code satisfies a specification but offers no guarantees about the specification's quality. This vulnerability allows models to pass verification using trivial or flawed specifications (Yan et al., 2025). While Yan et al. (2025) attempt to address this by comparing generated specifications against ground-truth specifications using a partial order, their method remains dependent on the quality and availability of trusted ground-truth. In contrast, our work introduces a formal equivalence metric that verifies the bidirectional correspondence between code and specification. This approach ensures the specification fully captures the program's behavior without relying on a ground-truth specification.

The absence of such a metric has hampered the creation of high-quality benchmarks for **autoformalization**. Existing datasets often lack the tripartite alignment of natural language, code, and formal specifications necessary for training (Lohn & Welleck, 2024; Loughridge et al., 2025; Dougherty & Mehta, 2025; Yan et al., 2025) or are small-scale due to the high cost of manual annotation (Misu et al., 2024; Sun et al., 2024; Miranda et al., 2025; Ye et al., 2025b). Attempts to automate equivalence checking have proven unreliable; for instance, Clover (Sun et al., 2024) relies on LLM-based judgments that suffer from high error rates. Addressing these deficiencies, we present **VeriEquiv-Bench**, a benchmark an order of magnitude larger than prior work. Enabled by our robust equivalence metric, it provides a large-scale, trustworthy resource for developing and evaluating models for verifiable code generation.

6 Conclusion

In this paper, we confront a foundational challenge hindering the development of reliable verifiable systems: the dependence on small, manually-annotated benchmarks for formal verification. This issue limits the scale and complexity of evaluation and has also introduced a ceiling by human knowledge. Our work breaks the dependency and introduces VeriEquivBench, a large-scale end-to-end formally verifiable code generation benchmark. Our novel automated equivalence score provides a rigorous evaluation signal without any need for human-written, ground-truth specifications. Second, our structured tagging system enables the scalable, automated synthesis of novel and complex problems, directly addressing the data generation bottleneck. By using VeriEquivBench to evaluate state-of-the-art LLMs, we have demonstrated that end-to-end verifiable code generation remains an open challenge, a fact obscured by the inflated success rates on simpler, older benchmarks. Following the recent discussions on self-evolving agents, our benchmark provides a scalable data generation engine and a reliable auto-evaluation metric, setting the groundwork to foster trustworthy AI agents with exact solution generation and sustainably supervise super-intelligence agents.

7 REPRODUCIBILITY STATEMENT

The code and our dataset are included in the supplementary material and will be publicly available after the double-blind review process for reproducibility.

8 ETHICS STATEMENT

This work does not present any foreseeable ethical concerns. The research involves only publicly available datasets and does not use or analyze sensitive or personally identifiable information.

REFERENCES

- Luogu online judge, 2025. URL https://www.luogu.com.cn/. Accessed: 2025-05-28.
- 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.
- Francois Chollet, Mike Knoop, Gregory Kamradt, Bryan Landers, and Henry Pinkard. Arc-agi-2: A new challenge for frontier ai reasoning systems. *arXiv preprint arXiv:2505.11831*, 2025.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.
- Domenico Cotroneo, Cristina Improta, Pietro Liguori, and Roberto Natella. Vulnerabilities in ai code generators: Exploring targeted data poisoning attacks. In *Proceedings of the 32nd IEEE/ACM International Conference on Program Comprehension*, pp. 280–292, 2024.
- David Dalrymple, Joar Skalse, Yoshua Bengio, Stuart Russell, Max Tegmark, Sanjit Seshia, Steve Omohundro, Christian Szegedy, Ben Goldhaber, Nora Ammann, et al. Towards guaranteed safe ai: A framework for ensuring robust and reliable ai systems. *arXiv preprint arXiv:2405.06624*, 2024.
- Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *International conference* on Tools and Algorithms for the Construction and Analysis of Systems, pp. 337–340. Springer, 2008.
- Leonardo De Moura, Soonho Kong, Jeremy Avigad, Floris Van Doorn, and Jakob Von Raumer. The lean theorem prover (system description). In *International Conference on Automated Deduction*, pp. 378–388. Springer, 2015.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Tianyu Liu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. A survey on incontext learning. 2023. doi: 10.48550/ARXIV.2301.00234. URL https://arxiv.org/abs/2301.00234.
- Quinn Dougherty and Ronak Mehta. Proving the coding interview: A benchmark for formally verified code generation. In 2025 IEEE/ACM International Workshop on Large Language Models for Code (LLM4Code), pp. 72–79. IEEE, 2025.
- Jiazhan Feng, Shijue Huang, Xingwei Qu, Ge Zhang, Yujia Qin, Baoquan Zhong, Chengquan Jiang, Jinxin Chi, and Wanjun Zhong. Retool: Reinforcement learning for strategic tool use in llms, 2025. URL https://arxiv.org/abs/2504.11536.
- András György, Tor Lattimore, Nevena Lazić, and Csaba Szepesvári. Beyond statistical learning: Exact learning is essential for general intelligence. *arXiv preprint arXiv:2506.23908*, 2025.

- Xuhan Huang, Qingning Shen, Yan Hu, Anningzhe Gao, and Benyou Wang. LLMs for mathematical modeling: Towards bridging the gap between natural and mathematical languages. In Luis Chiruzzo, Alan Ritter, and Lu Wang (eds.), *Findings of the Association for Computational Linguistics: NAACL 2025*, pp. 2678–2710, Albuquerque, New Mexico, April 2025. Association for Computational Linguistics. ISBN 979-8-89176-195-7. doi: 10.18653/v1/2025.findings-naacl. 146. URL https://aclanthology.org/2025.findings-naacl.146/.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*, 2024.
- Albert Qiaochu Jiang, Sean Welleck, Jin Peng Zhou, Timothee Lacroix, Jiacheng Liu, Wenda Li, Mateja Jamnik, Guillaume Lample, and Yuhuai Wu. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. In *The Eleventh International Conference on Learning Representations*.
- Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? In *12th International Conference on Learning Representations, ICLR 2024*, 2024.
- K Rustan M Leino. Dafny: An automatic program verifier for functional correctness. In *International conference on logic for programming artificial intelligence and reasoning*, pp. 348–370. Springer, 2010.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. Competition-level code generation with alphacode. *Science*, 378(6624):1092–1097, 2022.
- Evan Lohn and Sean Welleck. minicodeprops: a minimal benchmark for proving code properties. *arXiv preprint arXiv:2406.11915*, 2024.
- Chloe R Loughridge, Qinyi Sun, Seth Ahrenbach, Federico Cassano, Chuyue Sun, Ying Sheng, Anish Mudide, Md Rakib Hossain Misu, Nada Amin, and Max Tegmark. Dafnybench: A benchmark for formal software verification. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL https://openreview.net/forum?id=yBgTVWccIx.
- Thomas J. McCabe. A complexity measure. *IEEE Transactions on Software Engineering*, SE-2(4): 308–320, 1976. doi: 10.1109/TSE.1976.233837.
- Brando Miranda, Zhanke Zhou, Allen Nie, Elyas Obbad, Leni Aniva, Kai Fronsdal, Weston Kirk, Dilara Soylu, Andrea Yu, Ying Li, et al. Veribench: End-to-end formal verification benchmark for ai code generation in lean 4. In 2nd AI for Math Workshop@ ICML 2025, 2025.
- Md Rakib Hossain Misu, Cristina V Lopes, Iris Ma, and James Noble. Towards ai-assisted synthesis of verified dafny methods. *Proceedings of the ACM on Software Engineering*, 1(FSE):812–835, 2024.
- Martin Riddell, Ansong Ni, and Arman Cohan. Quantifying contamination in evaluating code generation capabilities of language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14116–14137, 2024.
- Parshin Shojaee, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar. The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity. *arXiv preprint arXiv:2506.06941*, 2025.
- David Silver, Satinder Singh, Doina Precup, and Richard S Sutton. Reward is enough. *Artificial Intelligence*, 299:103535, 2021.
- Chuyue Sun, Ying Sheng, Oded Padon, and Clark Barrett. Clover: Clo sed-loop ver ifiable code generation. In *International Symposium on AI Verification*, pp. 134–155. Springer, 2024.
- 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. *arXiv* preprint arXiv:2505.13938, 2025.

- Shangqing Tu, Kejian Zhu, Yushi Bai, Zijun Yao, Lei Hou, and Juanzi Li. Dice: Detecting in-distribution contamination in llm's fine-tuning phase for math reasoning. *arXiv* preprint *arXiv*:2406.04197, 2024.
 - Yutong Wang, Pengliang Ji, Chaoqun Yang, Kaixin Li, Ming Hu, Jiaoyang Li, and Guillaume Sartoretti. Mcts-judge: Test-time scaling in llm-as-a-judge for code correctness evaluation. *CoRR*, 2025a.
 - Zhijie Wang, Zijie Zhou, Da Song, Yuheng Huang, Shengmai Chen, Lei Ma, and Tianyi Zhang. Towards understanding the characteristics of code generation errors made by large language models. In 2025 IEEE/ACM 47th International Conference on Software Engineering (ICSE), pp. 717–717. IEEE Computer Society, 2025b.
 - Zihan Wang, Jiaze Chen, Zhicheng Liu, Markus Mak, Yidi Du, Geonsik Moon, Luoqi Xu, Aaron Tua, Kunshuo Peng, Jiayi Lu, Mingfei Xia, Boqian Zou, Chenyang Ran, Guang Tian, Shoutai Zhu, Yeheng Duan, Zhenghui Kang, Zhenxing Lin, Shangshu Li, Qiang Luo, Qingshen Long, Zhiyong Chen, Yihan Xiao, Yurong Wu, Daoguang Zan, Yuyi Fu, Mingxuan Wang, and Ming Ding. Aethercode: Evaluating llms' ability to win in premier programming competitions, 2025c. URL https://arxiv.org/abs/2508.16402.
 - 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, 2025. URL https://arxiv.org/abs/2504.14655.
 - Zhangchen Xu, Yang Liu, Yueqin Yin, Mingyuan Zhou, and Radha Poovendran. Kodcode: A diverse, challenging, and verifiable synthetic dataset for coding, 2025. URL https://arxiv.org/abs/2503.02951.
 - Chuanhao Yan, Fengdi Che, Xuhan Huang, Xu Xu, Xin Li, Yizhi Li, Xingwei Qu, Jingzhe Shi, Zhuangzhuang He, Chenghua Lin, et al. Re: Form–reducing human priors in scalable formal software verification with rl in llms: A preliminary study on dafny. *arXiv preprint arXiv:2507.16331*, 2025.
 - Tian Ye, Zicheng Xu, Yuanzhi Li, and Zeyuan Allen-Zhu. Physics of Language Models: Part 2.1, Grade-School Math and the Hidden Reasoning Process. In *Proceedings of the 13th International Conference on Learning Representations*, ICLR '25, April 2025a. Full version available at https://ssrn.com/abstract=5250629.
 - Zhe Ye, Zhengxu Yan, Jingxuan He, Timothe Kasriel, Kaiyu Yang, and Dawn Song. Verina: Benchmarking verifiable code generation. *arXiv preprint arXiv:2505.23135*, 2025b.
 - Boxi Yu, Yuxuan Zhu, Pinjia He, and Daniel Kang. Utboost: Rigorous evaluation of coding agents on swe-bench. *arXiv preprint arXiv:2506.09289*, 2025.
 - Andrew Zhao, Yiran Wu, Yang Yue, Tong Wu, Quentin Xu, Matthieu Lin, Shenzhi Wang, Qingyun Wu, Zilong Zheng, and Gao Huang. Absolute zero: Reinforced self-play reasoning with zero data. *arXiv preprint arXiv:2505.03335*, 2025.

A DETAILS ABOUT ALGORITHM TAGS

To assemble a suitable tag vocabulary, we first collect high-quality, high-frequency labels from Luogu—a competitive-programming platform with millions of users and an unusually fine-grained tag taxonomy—and treat them as a seed set. For each LeetCode problem, the model is prompted to pick the most relevant domain, data-structure, and algorithm tags from this pool, and is allowed to introduce new tags only when no suitable match exists. All model-selected tags are pooled, automatically partitioned into the three coarse categories, and then manually filtered in a single pass: hallucinated labels are removed, near-duplicates merged, and overly broad or overly narrow tags discarded. The resulting inventory contains over 500 clean triples that serve as the controlled vocabulary for subsequent tag-composition.



Figure 7: Statistics of algorithm tags

Tag category	Numbers	
Domain	53	
Data Structure	68	
Algorithm	480	

Figure 6: The fifteen most frequently used tags in our dataset.

The complete curated tag set is listed below, grouped under the three top-level categories: domain, data-structure, and algorithm.

Table 5: Domain tags

Category	Tags
Domain	Mathematics, Number Theory, Probability Theory, Combinatorial Mathematics, Linear Algebra, Computational Geometry, Plane geometry, Threedimensional computational geometry, Graph Theory, Simple Graph Theory, Game Theory, Information Theory, Dynamic Connectivity, expectation, Set Cover Problem, allocation problem, Extremum problem, path problem, Chess Board Problem, Stock Problem, Island Problem, Maze Problem, Josephus problem, Frobenius problem, N-Queens Problem, Knight's Tour Problem, Two-dimensional partial order problem, matching problem, Pairing problem, Interval problems, Knapsack problem, Subset Sum Problem, Jump Game, Maximum Subarray Problem, Maximum Subsequence Problem, Largest Rectangle in Histogram, longest chain, Path counting, Path Statistics, Connectivity, Reachability analysis, periodic, Discrete Event Simulation, Time constraint, Permutations and Combinations, Counting Principles, Inclusion-Exclusion Principle, Pigeonhole principle, Catalan number, Stirling numbers of the second kind, Combinatorial counting, Combinatorial Optimization, Mathematical Techniques

Table 6: Data Structure tags

Category	Tags
Data Structure	array, Two-dimensional array, Multidimensional array, sorted array, Circular array, tagged array, Difference Array, rolling array, Linked List, doubly linked list, Circular Linked List, Queue, deque, Priority Queue, Stack, monotonic stack, monotonic queue, tree, undirected tree, unrooted tree, Ring tree, Binary Tree, Complete Binary Tree, Perfect Binary Tree, Balanced Binary Tree, Binary Search Tree, Tree data structure, Trie, Segment Tree, Binary Indexed Tree, Heap, heap - min heap, Huffman tree, Set, Hash Table, Adjacency List, Adjacency Matrix, weight graph, Bipartite graph, Complete graph, Undirected graph, directed graph, Reverse graph, Star graph, Directed Acyclic Graph (DAG), Balanced tree, sparse matrix, Disjoint Set Union (DSU), Red-Black Tree, AVL Tree, B-Tree, B+ Tree, Skip List, Bloom Filter, LRU Cache, Prefix Tree, Suffix Tree, Suffix Array, Cartesian Tree, Splay Tree, Scapegoat Tree, Persistent Data Structure, Linear List, Sparse Table, Mo's Algorithm Structure, Leftist Tree, Fibonacci Heap, Pairing Heap

Tags

756 757

Category

Algorithm-1

Table 7: Algorithm tags

772

780

795

796

797

803

804

6
Compression algorithm, Dynamic Programming, Dynamic Programming - Linear DP, Dynamic Programming-LIS, Dynamic Programming-Prefix
Sum, Dynamic Programming - 0/1 Knapsack, Dynamic Programming - State
Compression, Dynamic Programming - Interval DP, Dynamic Programming
- 2D DP,Dynamic Programming - Prefix Sum Optimization,Dynamic Pro-
gramming - Top-Down, Dynamic Programming - Iterative, Dynamic Program-
ming, Compression algorithm, Dynamic Programming, Dynamic Programming
- Linear DP, Dynamic Programming-LIS, Dynamic Programming-Prefix
Sum, Dynamic Programming - 0/1 Knapsack, Dynamic Programming - State
Compression, Dynamic Programming - Interval DP, Dynamic Program-
ming - 2D DP,Dynamic Programming - Prefix Sum Optimization,Dynamic
Programming - Top-Down, Dynamic Programming - Iterative, Dynamic Pro-
gramming, State Compression DP, Dynamic Programming - Mathematical
Optimization, Digital DP, Count DP, Tree DP, knapsack DP, State Compression
DP,Dynamic Programming (DP),2D DP,Bidirectional DP,Sequence DP,Matrix
DP,State Machine DP,Bottom-up Dynamic Programming,Bidirectional
BFS,Multi-source BFS,0-1 BFS,Depth-First Search (DFS),Breadth-
First Search (BFS), Memoization, State space search, Heuristic search, state
search, Grid search, Path Finding, Binary search, Binary Search - Answer, Binary
Search - Right Boundary, Binary Search - Left Boundary, Binary Search -
Count, Binary Search - Peak Finding, Binary Search - Maximum Value, Binary
Search-Prefix Sum, Binary Search - Middle Element, Binary Search - Line
Search

Table 8: Algorithm tags

Tags

Sorting, Merge sort, Quick Sort, Three-way quicksort, Insertion Sort, Counting Sort, Bucket Sort, Sort-Custom Sort, Sorting Stable Sort, Sorting Lexicographical Order, Difference Sorting, multi-condition sorting, Wiggle Sort, in-place sorting, Topological sorting, Quick Select, KMP algorithm, Rabin-Karp algorithm, Manacher's algorithm, suffix array, suffix tree,Z-function,prefix function, string pattern matching, string card matching, backtracking, Enumeration, Binary Enumeration, Subset Enumeration, Combinatorial Enumeration, Two-dimensional ation, Simulation, Greedy, Greedy - Interval Operation, Divide quer, Divide and Conquer - String Splitting, Divide and Conquer - Closest Pair of Points in a Plane, Central Expansion Method, Staining method, Contribution method, sliding window, Two Pointers, Two Pointers - Sliding Window, Fast and slow pointers, Three Pointers, path compression, Path Tracing, Path reconstruction, Path Planning, Single-Source Shortest Path, Multi-Source Shortest Path, Second shortest circuit, Constrained Shortest Path, shortest path, Heapoptimized Dijkstra, Dijkstra's algorithm, Dijkstra's Algorithm Variant, Bellman-Ford algorithm, Floyd's cycle-finding algorithm, Kruskal's algorithm, Prim's algorithm, Minimum Spanning Tree, Bipartite Matching, Maximum Matching in Bipartite Graphs, Hungarian algorithm, Minimum Cost Maximum Flow, Graham scan, Welzl's algorithm, linear sieve, Euler sieve, Eratosthenes sieve, Prime Sieve, Euclidean algorithm, Bézout's identity, Bézout's theorem, Greatest Common Divisor (GCD), Least Common Multiple (LCM), Prime Number Check

Algorithm-2

Category

859

861 862

Table 9: Algorithm tags

811 812 Category **Tags** 813 Euclidean algorithm, Bézout's identity, Bézout's theorem, Greatest Common 814 Divisor(GCD), Least Common Multiple(LCM), Prime Number Check, Prime 815 factorization, Factorization, Integer factorization, Cantor expansion, Fast 816 exponentiation, Matrix Fast Exponentiation, Matrix multiplication, matrix 817 rotation, matrix transposition, Matrix operations, rotation matrix, flood fill 818 algorithm, A* algorithm, Tarjan's algorithm, Morris traversal, Preorder Traver-819 sal, Inorder Traversal, Postorder traversal, Level order traversal, Level Order Traversal, Reverse inorder traversal, zigzag traversal, spiral, traversal, Vertical 820 traversal, Vertical Order Traversal, Boundary traversal, Diagonal Traver-821 sal,2D matrix traversal,Traversal of 2D Array, Graph traversal,Linked 822 list traversal, Tree traversal, Directional traversal, Bidirectional 823 sal,reverse traversal,Reverse traversal,One-pass traversal,Path 824 tion, Path counting, Path Statistics, Path Construction, lexicographical com-825 parison,Lexicographically smallest path,Maximum Value Search,Maximum Value Maintenance, Range Maximum, Maximum Column Value, prefix 827 maximum, suffix minimum, suffix product, prefix product, Prefix Sum, Prefix 828 Sum - Difference, Prefix Sum - Modular Arithmetic, Prefix Sum - Bi-829 nary Search Optimization,2D prefix sum,suffix sum,partial sum, subarray 830 sum, submatrix sum, Area Sum, Area Calculation, ASCII code manipulation, Character Mapping, Character Count, character frequency, Digital encod-831 ing, Digital Parsing, Data Extraction, Number Reversal, Integer Reversal, Integer 832 Square Root, Integer Division, Fraction Addition and Subtraction, Fractional 833 Arithmetic,Fraction simplification, Score Calculation, percentile, Circular 834 shift, Loop Detection, Ring Detection, Periodic Assessment, Bracket Match-Algorithm-3 835 ing, Isomorphic Strings, String comparison, String Case Conversion, String 836 concatenation, string concatenation, String manipulation, String search, string 837 matching, String-Substring Comparison, string-replacement, String replace-838 ment, String trimming, string slicing, string splitting, String compression, String 839 decoding, string parsing, string continuity, substring matching, prefix match-840 ing, Prefix Check, Longest Common Prefix, Longest Common Suffix, Longest 841 Common Subsequence, Longest Common Subarray, Longest Repeating Substring, Longest Palindromic Subsequence, Longest Non-decreasing Sub-842 array, Longest Consecutive Sequence, longest consecutive characters, Word 843 Chain, Zigzag Conversion, palindrome, Expression parsing, Expression Evalua-844 tion, Reverse Polish Notation, Postfix expression, Operator precedence, Lexical 845 Analysis, parsing, Serialization, Deserialization, Encoding, decoding, Run-846 encoding,Set Operations,Set Intersection, Bitwise 847 tion, Bitwise operation optimization, Bitwise Operations - State Compres-848 sion, bitmask, Bitwise OR, AND operation, XOR, binary, Binary Addition, binary 849 splitting, Binary counting, bit count, Hamming distance, Two's 850 ment.Modular arithmetic, modulo 3 operation, Congruence, Congruence 851 theorem, divisible, Divisibility property, divisor, perfect square, square num-852 ber, Perfect number, Ugly number, trailing zeros, digit separation, Digital Processing, Digital Sum, Gray code, Permutation, Next Permuta-853 tion, Arrangement, Permutation ring, Cyclic permutation, Pascal's 854 gle, Fermat's theorem on sums of two, squares, Pythagorean theorem, Triangle 855 inequality, absolute value, absolute value inequality, Big Integer Addition, High 856 precision 858

Table 10: Algorithm tags

865 866 Category **Tags** 867 Floating-point processing, Floating-point comparison, floating-point 868 precision,Linear equation,polynomial,Complex Number Operarepresentation, recurring number tions,Rational decimal,factorial,Sum 870 Squares,Sum,Summation formula, arithmetic sequence, Arithmetic 871 sequence summation,path sum,Maximum Sum Path,Maximum spac-872 ing, Neighbor Count, Adjacent elements, Adjacent Element Difference, Global 873 Inversion, Local inversion pairs, Inversion pair, anagram, vowel substitution,coordinate,2D coordinates,coordinate system,coordinate 874 son, coordinate translation, coordinate compression, 2D offset, 2D plane, 3D 875 space, collinear points, Collinearity detection, convex hull, minimum bound-876 ing rectangle, Triangle Area, Rectangle Area Calculation, Overlapping 877 Calculation, Rectangle Intersection, Circle-Rectangle Area Intersection 878 Detection, Minimum Enclosing Circle, Spatial segmentation,2D 879 ting, Spatial optimization, Space complexity optimization, Constant space 880 complexity, Linear space complexity, Time complexity analysis, Linear time complexity,Linear scan,Pruning,Preprocessing,preprocessing,Offline 882 update,Dynamic processing, Dynamic Maintenance, Dynamic 883 Interval, Dynamic Range Maintenance, Single-point nance modifica-884 tion,Range query,Interval computation,Interval Statistics, Range date, Interval Merging, Interval coverage, Interval Scheduling, Range 885 extrema, Path Intersection Detection, Distance calculation, Euclidean tance, Manhattan distance, Chebyshev distance, projection, cross 887 uct, Polar sorting, construct, Binary Construction, Tree Construction, Tree Reconstruction, Sequence Reconstruction, Constructing the answer in Algorithm-4 889 reverse order,reverse,Reverse Linked List,Linked List Reversal,String 890 Reversal, Array Rearrangement, Linked List Reordering, Node switch-891 ing, Segmentation, Split Array, split string, Split and Merge, Convert 1D 892 Array to 2D Array,matrix,2D matrix,sparse matrix,ordered matrix,Rectangle 893 Coverage, Adjacency Matrix, Tree deletion operation, Tree depth, Tree Cen-894 troid, Tree Diameter, subtree, Subtree Sum, leaf node, intermediate node, dummy 895 node, sentinel node, Middle of the Linked List, indegree, indegree and outdegree,degree sequence,Monotonicity,Monotonicity Check,monotonic 896 array, Decision Monotonicity, Symmetric, Boolean operations, Logical Oper-897 ations, Conditional statement, Filter Criteria, Polarity, Parity Check, Boundary check.Boundary handling,Edge case handling,Status Check.Status 899 Log, State transition, State Machine, Finite State Automaton, Priority, han-900 dling, Query Processing, Path processing, Overflow handling, Carry han-901 dling, Recursion, recursive, Inductive method, derivation, traverse, Array traver-902 sal, Grid traversal, directional search, State compression, Handling Duplicate 903 Elements, deduplication, Enumeration optimization, Sequence 904 son, comparison function, Comparator, Regular Expression, Pointer manipulation, Method chaining, Swap operation, Displacement operation, Row and Col-905 906 umn Operations, product, Multiplication Principle, Exponentiation, Base, Base Conversion, Clock issues, loop section, IP address, reordering, Partial Or-907 dering, Equation Solving, Randomization, reverse thinking, Horse Racing 908 Strategy, Connected component, Connected Component, Strongly Connected 909 Component, Lowest Common Ancestor (LCA), Eulerian circuit, Hamiltonian 910 path 911 912

B PIPELINE OF TAG COMPOSITION

Figure 8 illustrates our pipeline for generating new programming problems through tag composition. The process begins by creating a candidate pool of 36 tags, randomly selecting 12 from each of our three categories: domain, algorithm, and data structure. This pool is provided to an LLM, which is prompted to select a coherent subset of three to eight tags that form a promising basis for a new problem. Using this selected combination, we then instruct the LLM to generate a complete task, comprising a problem description, corresponding unit tests, and a Python solution. As a final quality control step, we filter these generations by executing the unit tests. We retain only those instances where the generated Python code passes all tests, ultimately yielding a dataset of 300 validated programs.

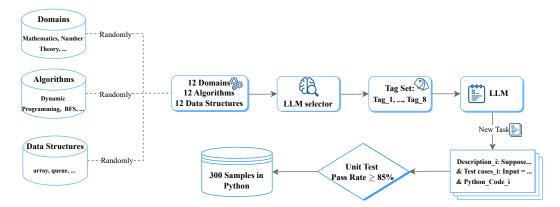


Figure 8: The pipeline for the tag compositon process.

C PROMPT TEMPLATES

[Task]

You have three categories of tags: domain, algorithm, and data_structure, each containing 12 tags. Your task is to select a combination of 3–8 tags from these categories to form a coherent programming problem. The problem should have a specified difficulty level: easy, medium, or hard. Ensure the selected tags are compatible and can logically form a single problem. Provide the chosen tags, the difficulty level.

[Requirements]

- 1. The task is clearly defined, specifying the need to select 3–8 tags from three categories (domain, algorithm, data_structure) to form a coherent programming problem with a specified difficulty level.
- 2.Requirements outline the tag selection process, ensuring compatibility and a reasonable tag collection, the need for a difficulty level.

 3.The selected tags must be compatible and form a reasonable tag collection that results in a practical and solvable programming problem.
- 4.The problem must be assigned one of three difficulty levels: easy, medium, or hard, reflecting the complexity of the problem based on the selected tags.

Figure 9: The prompt is for selecting useful tags. We feed the model the 36 real tags from 3 categories randomly that will later drive new-problem generation, it returns the 3–8 tags that form the most promising combination.

```
1026
           [Task]
1027
           You are an expert algorithm problem creator. Your task is to create an easy or medium difficulty ranking original coding problem using
1028
          the given algorithm tags. Analyze the given tags to generate a new problem. The problem should be completely original coding problem
1029
          that is NOT from any existing platforms (LeetCode, Codeforces, etc.) or textbooks.
1030
           [Requirements]
1031
          1. Create a truly novel problem scenario with unique constraints
1032
          2. Combine the given tags in innovative ways
1033
          3. Ensure the problem is solvable but challenging
          4. Provide clear problem statement, examples, and constraints
1034
          5. Rate the difficulty(easy, medium, hard) appropriately
1035
           [Algorithm tags]
1036
1037
           Output Format
1038
          This is the ouput format, You must respond in this specified output format:
1039
           <--Problem Begin-->
1040
           json
1041
           <--Problem End-->
1042
           [Example]
1043
          Below is an example output format.
           <--Problem Begin-->
1044
1045
1046
               "title": "Find Minimum in Rotated Sorted Array",
1047
               "difficulty": "medium",
               "description": "Suppose an array of length n sorted in ascending
1048
                    order is rotated between 1 and n times. For example, the array
1049
                     nums = [0,1,2,4,5,6,7] might become [4,5,6,7,0,1,2] if it was
1050
                     rotated 4 times. Given the sorted rotated array nums of
1051
                    unique elements, return the minimum element of this array. You
1052
                     must write an algorithm that runs in O(log n) time.",
                "input_format": "An integer array nums of length n where 1 <= n
1053
                    <= 5000",
1054
               "output_format": "Return the minimum value in the array",
1055
               "constraints": [
1056
                     "1 <= nums.length <= 5000",
                     "-5000 \le nums[i] \le 5000",
1057
                     "All the integers of nums are unique",
1058
                     "nums is sorted and rotated between 1 and n times"
                "examples": [
                          "input": "[3,4,5,1,2]",
1062
                          "output": "1",
1063
                          "explanation": "The original array was [1,2,3,4,5]
1064
                               rotated 3 times."
1065
1067
                          "input": "[4,5,6,7,0,1,2]",
                          "output": "0"
1068
                          "explanation": "The original array was [0,1,2,4,5,6,7]
1069
                               and it was rotated 4 times."
1070
1071
               "tags_used": ["Array", "Binary Search"],
               "time_complexity": "O(log n)",
               "space_complexity": "O(1)'
1074
1075
1076
           <--Problem End-->
```

Figure 10: The prompt uses the previously obtained real tags to generate a brand-new problem together with its unit tests.

D EXPERIMENTAL RESULT DETAILS

D.1 AUTOFORMALIZATION-EQUIVALENCE CHECK PASSING RATE

As shown in Figure 2 Pipeline 2, the autoformalized specifications are fed into an equivalence check with the corresponding NL query. The step of re-generating Python code and testing the code on unit tests filters out 13.32% of mismatched specifications. We do not further filter out data according to the LLM-as-a-judge results, to provide a fair evaluation on how effective this validation metric can be.

Table 11: The table reports the filter rate for autoformalization.

Dataset	Average Number of Unit Tests	Passing Rate
LeetCode-Train	89.15	86.68%

D.2 QUALITY ASSURANCE OF PYTHON2DAFNY TRANSFORMATION

This section lists the question ID randomly chosen to test the code transformation quality, matching the one provided in the original LeetCode dataset.

Table 12: Our Python to Dafny transformation step gives high-quality data with 90% passing all unit tests.

Q2914: 116/116 tests passed	Q3099: 53/53 tests passed	Q1419: 197/197 tests passed
Q118: 16/16 tests passed	Q1278: 89/89 tests passed	Q2549: 50/50 tests passed
Q767: 11/56 tests passed	Q276: 103/103 tests passed	Q3019: 78/78 tests passed
Q1910: 130/130 tests passed	Q2414: 115/115 tests passed	Q1247: 103/103 tests passed
Q2288: 106/108 tests passed	Q1358: 130/130 tests passed	Q3183: 56/56 tests passed
Q1155: 108/108 tests passed	Q1698: 84/84 tests passed	Q2582: 94/94 tests passed
Q1931: 61/61 tests passed	Q2062: 110/110 tests passed	_

D.3 TRAINING CURVES ON AUXILIARY TASKS

We use the 14B SFT model provided by the Veri-Code Team and their code to RL-train models using GRPO.

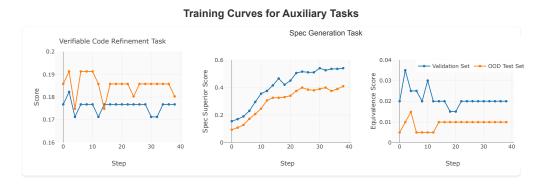


Figure 11: During the verifiable code refinement task, the model barely improves, demonstrating that RL training is not enough. During the spec generation task, the generated specification quality keeps enhancing, but still fails to capture code behaviours without ambiguities.

D.4 DETAILS ABOUT DAFNYSYNTHESIS INSPECTION

This section details our analysis of 14 ground-truth samples identified as problematic. Our investigation revealed that two samples failed initial verification due to implementation errors or timeouts,

precluding further analysis. These were #566 and #632, the latter of which was previously reported by Clover (Sun et al., 2024).

The primary issue in the remaining 12 samples was specification ambiguity stemming from insufficient post-conditions. We successfully rectified this in eight cases by strengthening their post-conditions, with the fixes validated through equivalence testing. Although our refinements improved the specifications for two other samples, they still did not pass the equivalence check. We were unable to resolve the ambiguities in the final two samples.

A significant portion of these ambiguous samples were newly discovered. Specifically, eight samples (#579, #602, #625, #629, #733, #755, #793, #807) were not documented in the prior work by Clover (Sun et al., 2024). Of these, we successfully fixed five (#625, #733, #755, #793, #807) and refined one (#602). Corresponding code examples are shown in Figures 12-19.

Regarding the issues previously reported by Clover, our findings for samples #567, #576, #644, and #803 largely concur. We fixed three (#567, #644, #803) and refined one (#576), with details in Figures 20-22. Conversely, sample #472 passed our equivalence checks, which aligns with Clover's classification of its issue as a natural-language inconsistency rather than a specification defect. As noted, sample #632 was excluded due to a timeout.

```
1193
1194
       predicate InArray(a: array<int>, x: int)
1195
           reads a
1196
1197
           exists i :: 0 \le i \le a.Length && a[i] == x
1198
1199
       method DissimilarElements(a: array<int>, b: array<int>) returns (result:
1200
           seq<int>)
1201
           ensures forall x :: x in result ==> (InArray(a, x) != InArray(b, x))
1202
           ensures forall i, j :: 0 \le i \le j \le |result| ==> result[i] != result[
1203
1204
           ####### The post-conditions here omit that the first half of result
               is in array a, while the second half is in b.
1205
1206
           var res: seq<int> := [];
1207
           for i := 0 to a.Length
1208
               invariant 0 <= i <= a.Length</pre>
1209
               invariant forall x :: x in res ==> InArray(a, x)
               invariant forall x :: x in res ==> InArray(a, x) != InArray(b, x)
1210
               invariant forall i, j :: 0 \le i \le j \le |res| \Longrightarrow res[i] != res[j]
1211
1212
               if !InArray(b, a[i]) && a[i] !in res
1213
               {
1214
                    res := res + [a[i]];
1215
1216
1217
           ghost var partialSize := |res|;
1218
           for i := 0 to b.Length
1219
               invariant 0 <= i <= b.Length</pre>
               invariant forall k :: partialSize <= k < |res| ==> InArray(b, res
1220
                    [k])
1221
               invariant forall k :: 0 \le k \le |res| ==> InArray(a, res[k]) !=
1222
                   InArray(b, res[k])
1223
               invariant forall i, j :: 0 <= i < j < |res| ==> res[i] != res[j]
1224
               if !InArray(a, b[i]) && b[i] !in res
1225
               {
1226
                    res := res + [b[i]];
1227
1228
           }
1229
           result := res;
1230
1231
1232
```

Figure 12: An unresolved specification ambiguity in DafnySynthesis sample #579. The post-condition is insufficient because it fails to enforce the preservation of the relative order of elements from the input array in the output.

1289

1290

```
1244
1245
1246
1247
       method FindFirstRepeatedChar(s: string) returns (found: bool, c: char)
1248
           ensures found ==> exists i, j :: 0 <= i < j < |s| && s[i] == s[j]
             && s[i] == c && (forall k, l :: 0 <= k < l < j
1249
             && s[k] == s[l] ==> k >= i)
1250
           ensures !found ==> (forall i, j :: 0 \le i \le j \le |s| ==> s[i] != s[j])
1251
           ######## The added post-condition
1252
           ensures !found ==> c == ' '
1253
           #######
1254
           c := ' ';
1255
           found := false;
1256
           var inner_found := false;
1257
           var i := 0;
1258
           while i < |s| \&\& !found
1259
               invariant 0 <= i <= |s|
               invariant found == inner_found
1260
1261
               invariant found ==> exists ii, jj :: 0 <= ii < i</pre>
1262
                  && ii < jj < |s| && s[ii] == s[jj] && s[ii] == c
1263
                  && (forall k, l :: 0 <= k < l < jj && s[k] == s[l] ==> k >= ii)
1264
               invariant !found <==> (forall ii, jj :: 0 <= ii < i</pre>
1265
                 && ii < jj < |s| ==> s[ii] != s[jj])
1266
               \#\#\#\#\#\#\# The added loop invariant
1267
               invariant !found ==> c == ' '
1268
               #######
1269
               var j := i + 1;
1270
               while j < |s| && !inner_found
1271
                   invariant i < j <= |s|</pre>
1272
                    invariant inner_found ==> exists k :: i < k < |s|</pre>
1273
                      && s[i] == s[k] && s[i] == c
                    invariant !inner_found
1274
                      <=> (forall k :: i < k < j ==> s[i] != s[k])
1275
                    ####### 

The added loop invariant
1276
                    invariant !inner_found ==> c == ' '
1277
                    invariant !found
1278
                    #######
1279
                    if s[i] == s[j] {
1280
                        inner_found := true;
1281
                        c := s[i];
1282
                    }
1283
                    j := j + 1;
1284
               found := inner_found;
1285
               i := i + 1;
1286
1287
       }
1288
```

Figure 13: A refined but unfixed specification for sample #602. While the shown refinement fails the equivalence test, a stricter post-condition (k > 1) could not be verified due to a timeout.

```
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
       method SwapFirstAndLast(a: array<int>)
1316
           requires a.Length > 0
1317
           modifies a
           ######## The added post-condition
1318
           ensures a.Length == old(a.Length)
1319
           #######
1320
           ensures a[0] == old(a[a.Length - 1])
1321
           ensures a[a.Length - 1] == old(a[0])
1322
           ensures forall k :: 1 \le k \le a.Length - 1 ==> a[k] == old(a[k])
1323
           var tmp := a[0];
1324
           a[0] := a[a.Length - 1];
1325
           a[a.Length - 1] := tmp;
1326
      }
1327
1328
```

Figure 14: A successfully resolved specification ambiguity in DafnySynthesis sample #625. The original specification was ambiguous as it lacked a constraint on the output array's length. The ambiguity was rectified by introducing a post-condition ensuring the length remains invariant.

```
1354
1355
1356
1357
      predicate IsEven(n: int)
1358
1359
           n % 2 == 0
1360
1361
      method FindEvenNumbers(arr: array<int>) returns (evenList: seq<int>)
1362
1363
           ensures forall i :: 0 <= i < |evenList| ==> IsEven(evenList[i])
1364
           && evenList[i] in arr[..]
1365
           ensures forall i :: 0 <= i < arr.Length && IsEven(arr[i])</pre>
1366
           ==> arr[i] in evenList
           ####### The post-conditions here do not ensures the order preserving
1367
               between the input array and output array
1368
1369
           evenList := [];
1370
           for i := 0 to arr.Length
1371
               invariant 0 <= i <= arr.Length</pre>
               invariant 0 <= |evenList| <= i</pre>
1372
               invariant forall k :: 0 <= k < |evenList| ==> IsEven(evenList[k])
1373
                    && evenList[k] in arr[..]
1374
               invariant forall k :: 0 \le k \le i \&\& IsEven(arr[k]) ==> arr[k] in
1375
                   evenList
1376
               if IsEven(arr[i])
1377
1378
                   evenList := evenList + [arr[i]];
1379
1380
1381
1382
      method FindEvenNumbers_check(arr: array<int>) returns (evenList: seq<int
1383
          >)
1384
1385
         evenList := *;
1386
         assume forall i :: 0 <= i < |evenList| ==> IsEven(evenList[i]) &&
            evenList[i] in arr[..];
1387
         assume forall i :: 0 <= i < arr.Length && IsEven(arr[i]) ==> arr[i] in
1388
            evenList;
1389
         var val_0 :=FindEvenNumbers(arr);
1390
         assert evenList[..] == val_0[..];
1391
1392
1393
```

Figure 15: An unresolved specification ambiguity in DafnySynthesis sample #629. The post-condition is insufficient because it fails to enforce the preservation of the relative order of elements from the input array.

```
1411
1412
1413
1414
       method FindFirstOccurrence(arr: array<int>, target: int) returns (index:
1415
           requires arr != null
1416
           requires forall i, j :: 0 <= i < j < arr.Length ==> arr[i] <= arr[j]</pre>
1417
           ensures 0 <= index < arr.Length ==> arr[index] == target
1418
           ensures index == -1 ==> forall i :: 0 <= i < arr.Length ==> arr[i] !=
1419
1420
           ensures forall i :: 0 <= i < arr.Length ==> arr[i] == old(arr[i])
           \#\#\#\#\#\#\# The added post-condition
1421
           ensures 0 <= index < arr.Length || index == -1</pre>
1422
           ensures 0 <= index < arr.Length ==> ((forall i :: 0 <= i < index ==>
1423
               arr[i] < arr[index]) && (forall j :: index <= j < arr.Length ==>
1424
               arr[j] >= arr[index]))
1425
           #######
       {
1426
           index := -1;
1427
           for i := 0 to arr.Length
1428
               invariant 0 <= i <= arr.Length</pre>
1429
               invariant index == -1 ==> forall k :: 0 <= k < i ==> arr[k] !=
1430
               invariant 0 <= index < i ==> arr[index] == target
1431
               invariant forall k :: 0 \le k \le arr.Length ==> arr[k] == old(arr[k])
1432
                   ])
1433
1434
               if arr[i] == target
1435
                   index := i;
1436
                   break;
1437
1438
               if arr[i] > target
1439
               {
1440
                   break;
1441
1442
1443
```

Figure 16: A successfully resolved specification ambiguity in DafnySynthesis sample #733. The original specification was insufficient, lacking detail for cases where the input index is non-negative. The issue was fixed by refining the post-condition to explicitly define the expected behavior for this scenario.

```
1458
       function MinPair(s: seq<int>) : (r: int)
1459
           requires |s| == 2
1460
           ensures s[0] \le s[1] \le r = s[0]
1461
           ensures s[0] > s[1] ==> r == s[1]
1462
           if s[0] \le s[1] then s[0] else s[1]
1463
1464
       function min(s: seq<int>) : (r: int)
1465
           requires |s| >= 2
1466
           ensures forall i :: 0 \le i \le |s| \Longrightarrow r \le s[i]
1467
           if |s| == 2 then MinPair(s)
1468
           else MinPair([s[0], min(s[1..])])
1469
1470
       method SecondSmallest(s: array<int>) returns (secondSmallest: int)
1471
           requires s.Length >= 2
           requires exists i, j :: 0 <= i < s.Length && 0 <= j < s.Length
1472
             && i != j \&\& s[i] == min(s[..]) \&\& s[j] != s[i]
1473
           ensures exists i, j :: 0 <= i < s.Length && 0 <= j < s.Length</pre>
1474
             && i != j \&\& s[i] == min(s[..]) \&\& s[j] == secondSmallest
           ensures forall k :: 0 \le k \le s.Length && s[k] != min(s[..])
1476
             ==> s[k] >= secondSmallest
1477
           ####### 

The added post-condition
           ensures (exists i, j :: i != j && 0 <= i < s.Length
1478
             && 0 <= j < s.Length && s[i] == s[j] && s[i] == min(s[..]))
1479
             ==> secondSmallest == min(s[..])
1480
           ensures !(exists i, j :: i != j && 0 <= i < s.Length
1481
             && 0 \le j \le s.Length && s[i] == s[j] && s[i] == min(s[..]))
             ==> ( (exists k :: 0 \le k \le s.Length \&\& s[k] == secondSmallest)
1482
               && (forall k :: 0 \le k \le s.Length && s[k] > min(s[..])
1483
               ==> s[k] >= secondSmallest) && secondSmallest > min(s[..]) )
1484
           #######
1485
1486
           var minIndex := 0;
           var secondMinIndex := 1;
1487
           if s[1] < s[0] {
1488
               minIndex := 1;
1489
               secondMinIndex := 0;
1490
1491
           for i := 2 to s.Length
           invariant 0 <= i <= s.Length</pre>
1492
           invariant 0 <= minIndex < i</pre>
1493
           invariant 0 <= secondMinIndex < i</pre>
1494
           invariant minIndex != secondMinIndex
1495
           invariant forall k :: 0 \le k \le i ==> s[k] >= s[minIndex]
1496
           invariant forall k :: 0 \le k \le i \&\& k != minIndex ==> s[k] >= s[
               secondMinIndex]
1497
1498
               if s[i] < s[minIndex] {</pre>
1499
                    secondMinIndex := minIndex;
1500
                    minIndex := i;
               } else if s[i] < s[secondMinIndex] {</pre>
1501
1502
                    secondMinIndex := i;
1503
1504
1505
           secondSmallest := s[secondMinIndex];
1506
1507
```

Figure 17: A successfully resolved specification ambiguity in DafnySynthesis sample #755. The original specification was insufficient, failing to distinguish between cases with a unique minimum value and those with multiple occurrences of the minimum. The ambiguity was rectified by refining the post-condition to explicitly detail the expected behavior for both scenarios.

1510

```
1519
1520
1521
       method LastPosition(arr: array<int>, elem: int) returns (pos: int)
1522
           requires arr.Length > 0
1523
           requires forall i, j :: 0 <= i < j < arr.Length ==> arr[i] <= arr[j]</pre>
           ####### 

Original post-condition
1524
           // ensures pos == -1 || (0 <= pos < arr.Length && arr[pos] == elem &&
1525
                (pos \le arr.Length - 1 \mid | arr[pos + 1] > elem))
1526
           #######
1527
           ####### 

The fixed post-condition
1528
           ensures pos == -1 <==> (forall j :: 0 <= j < arr.Length ==> arr[j] !=
               elem)
1529
           ensures pos !=-1 <==> (0 <= pos < arr.Length && arr[pos] == elem &&
1530
               (pos == arr.Length - 1 \mid | arr[pos + 1] > elem))
1531
           #######
1532
           ensures forall i :: 0 <= i < arr.Length ==> arr[i] == old(arr[i])
1533
           pos := -1;
1534
           for i := 0 to arr.Length #### Originally, the upper bound is arr.
1535
               Length - 1, but it was buggy
1536
               invariant 0 <= i <= arr.Length</pre>
1537
               ####### 

####### 

Original loop invariant
               // invariant pos == -1 || (0 <= pos < i && arr[pos] == elem && (
1538
                   pos == i - 1 \mid\mid arr[pos + 1] > elem))
1539
               \#\#\#\#\#\#\# The fixed loop invariant
1540
               invariant pos == -1 <==> (forall j :: 0 <= j < i ==> arr[j] !=
1541
1542
               invariant pos != -1 <==> (0 <= pos < i && arr[pos] == elem && (
1543
                   pos == i - 1 || arr[pos + 1] > elem))
               #######
1544
               invariant forall k :: 0 <= k < arr.Length ==> arr[k] == old(arr[k]
1545
1546
1547
               if arr[i] == elem
1548
               {
                   pos := i;
1549
1550
           }
1551
1552
```

Figure 18: A successfully resolved specification ambiguity in DafnySynthesis sample #793. The original specification was insufficient as it failed to define distinct behaviors based on the sign of the input parameter 'pos'. The ambiguity was rectified by refining the post-condition to explicitly handle the cases where 'pos' is negative and non-negative, respectively.

```
1573
1574
1575
1576
1577
1578
       predicate IsOdd(x: int)
1579
           x % 2 != 0
1580
1581
1582
       method FindFirstOdd(a: array<int>) returns (found: bool, index: int)
1583
           requires a != null
1584
           ensures !found ==> forall i :: 0 <= i < a.Length ==> !IsOdd(a[i])
           ensures found ==> 0 <= index < a.Length && IsOdd(a[index])</pre>
1585
               && forall i :: 0 <= i < index ==> !IsOdd(a[i])
1586
           ######## The added post-condition
1587
           ensures !found ==> index == a.Length
1588
           #######
1589
           found := false;
1590
           index := 0;
1591
           while (index < a.Length)
1592
               invariant 0 <= index <= a.Length</pre>
1593
               invariant !found ==> forall i :: 0 <= i < index ==> !IsOdd(a[i])
1594
               invariant found ==> IsOdd(a[index - 1]) && forall i :: 0 <= i <</pre>
                   index - 1 ==> !IsOdd(a[i])
1595
1596
               if IsOdd(a[index])
1597
1598
                    found := true;
1599
                    return;
1600
               index := index + 1;
1601
1602
1603
1604
```

Figure 19: A successfully resolved specification ambiguity in DafnySynthesis sample #807. The original specification was insufficient, as it only described the behavior for successful outcomes. The ambiguity was resolved by strengthening the post-condition to explicitly define the program's state in failure cases, ensuring comprehensive and predictable behavior.

```
1627
1628
1629
1630
       method IsSorted(a: array<int>) returns (sorted: bool)
1631
           requires a.Length > 0
           ####### 

◆ Original post-condition
1632
           // ensures sorted <== forall i, j :: 0 <= i < j < a.Length</pre>
1633
              ==> a[i] <= a[j]
1634
           // ensures !sorted ==> exists i, j :: 0 <= i < j < a.Length
1635
             && a[i] > a[j]
1636
            #######
           \#\#\#\#\#\#\# The fixed post-condition
1637
           ensures sorted <==> forall i, j :: 0 <= i < j < a.Length</pre>
1638
             ==> a[i] <= a[j]
1639
           #######
1640
1641
           sorted := true;
           for i := 0 to a.Length - 1
1642
                invariant 0 <= i < a.Length</pre>
1643
                ####### \downarrow Original loop invariant
1644
                // invariant sorted \leftarrow forall k, l :: 0 \leftarrow k < l < i
1645
                  ==> a[k] <= a[l]
1646
                // invariant !sorted ==> exists k :: 0 \le k \le a[k] > a[k+1]
                #######
1647
                \#\#\#\#\#\#\# The fixed post-condition
1648
                invariant sorted <==> forall k, l :: 0 <= k < l <= i</pre>
1649
                  ==> a[k] <= a[l]
1650
                #######
1651
1652
                if a[i] > a[i + 1]
1653
                {
1654
                    sorted := false;
1655
                    break:
1656
1657
           sorted := sorted;
1658
1659
1660
```

Figure 20: A successfully resolved specification ambiguity in DafnySynthesis sample #567, an issue also identified by the Clover. The original post-condition was overly permissive, stating only a sufficient condition for the desired outcome. The ambiguity was rectified by strengthening this to a necessary and sufficient condition (an equivalence).

```
1678
1679
1680
1681
       method Reverse(a: array<int>)
1682
         modifies a
            \#\#\#\#\#\#\# The added post-condition
1683
         ensures a.Length == old(a.Length)
1684
            #######
1685
         ensures forall k :: 0 \le k \le a.Length == a[k] == old(a[(a.Length-1) -
1686
             k])
1687
         var l := a.Length - 1;
1688
         var i := 0;
1689
         while (i < l-i)
1690
            invariant 0 <= i <= (1+1)/2</pre>
1691
            invariant forall k :: 0 <= k < i \mid \mid 1-i < k <= 1 ==> a[k] == old(a[1-i])
1692
            invariant forall k :: i \le k \le l-i \Longrightarrow a[k] \Longrightarrow old(a[k])
1693
1694
           a[i], a[l-i] := a[l-i], a[i];
1695
            i := i + 1;
1696
1697
       method ReverseUptoK(s: array<int>, k: int)
1698
           modifies s
1699
            requires 2 <= k <= s.Length
1700
            ######## The added post-condition
1701
           ensures s.Length == old(s.Length)
1702
            #######
           ensures forall i :: 0 \le i \le k \Longrightarrow s[i] \Longrightarrow old(s[k-1-i])
1703
           ensures forall i :: k <= i < s.Length ==> s[i] == old(s[i])
1704
1705
         var 1 := k - 1;
1706
         var i := 0;
1707
         while (i < l-i)
            invariant 0 <= i <= (1+1)/2;
1708
            invariant forall p :: 0 <= p < i \mid \mid 1-i < p <= 1 ==> s[p] == old(s[1-i])
1709
               p]);
1710
            invariant forall p :: i \le p \le l-i ==> s[p] == old(s[p]);
1711
                invariant forall p :: k \le p \le s.Length ==> s[p] == old(s[p])
1712
           s[i], s[l-i] := s[l-i], s[i];
1713
            i := i + 1;
1714
1715
1716
```

Figure 21: A successfully resolved specification ambiguity in DafnySynthesis sample #644, an issue also identified by the Clover. The original specification was ambiguous as it lacked a constraint on the output array's length. The ambiguity was rectified by introducing a post-condition ensuring the length remains invariant.

```
1742
1743
1744
      method IsPerfectSquare(n: int) returns (result: bool)
           requires n >= 0
1745
           ####### \downarrow Original post-condition
1746
           // ensures result == true ==> (exists i: int :: 0 <= i <= n && i * i
1747
1748
           // ensures result == false ==> (forall a: int :: 0 < a*a < n ==> a*a
1749
              !=n)
           #######
1750
           ####### 

The fixed post-condition
1751
           ensures result <==> (exists i: int :: 0 <= i <= n && i * i == n)</pre>
1752
           #######
1753
           var i := 0;
1754
           while (i * i < n)
1755
               invariant 0 <= i <= n
1756
               invariant forall k :: 0 \le k \le i ==> k * k \le n
1757
1758
               i := i + 1;
1759
           }
           return i * i == n;
1760
1761
```

Figure 22: A successfully resolved specification ambiguity in DafnySynthesis sample #803, an issue also identified by the Clover. The original post-condition was overly permissive, stating only necessary conditions for the desired outcome. The ambiguity was rectified by strengthening this to a necessary and sufficient condition (an equivalence).

```
1789
1790
1791
1792
1793
       method IsSublist(sub: seq<int>, main: seq<int>) returns (result: bool)
           ####### 

Original post-condition
1794
           // ensures true <== (exists i :: 0 <= i <= |main| - |sub| && sub ==
1795
              main[i..i + |sub|])
1796
           #######
1797
           \#\#\#\#\#\#\# The refined post-condition
           ensures result ==> (exists i :: 0 <= i <= |main| - |sub| && sub ==</pre>
1798
              main[i..i + |sub|])
1799
           ensures result ==> (exists i :: |sub| <= i <= |main| && sub == main[i</pre>
1800
                - |sub|..i])
1801
           #######
1802
       {
           if |sub| > |main| {
1803
               return false;
1804
1805
           result := false;
1806
           for i := 0 to |main| - |sub| + 1
1807
               ####### 

The original loop invariant
               // invariant result ==> (exists j :: 0 <= j < i && sub == main[j</pre>
1808
                   ..j + |sub|])
1809
               #######
1810
               ######## The refined loop invariant
1811
               invariant 0 <= i <= |main| - |sub| + 1</pre>
1812
               #######
1813
1814
               if sub == main[i..i + |sub|] {
1815
                   result := true;
1816
1817
1818
           result := false;
1819
1820
```

Figure 23: An unresolved specification ambiguity in DafnySynthesis sample #576, an issue also identified by the Clover. The original post-condition was effectively meaningless, providing no meaningful constraints. Although the post-condition was refined to be more specific, the resulting specification still fails to pass the equivalence test, indicating that the ambiguity has not been fully resolved and requires further investigation.

E THE USE OF LARGE LANGUAGE MODELS

Multiple LLM products, including GPT-5 and Gemini-2.5-pro, are deployed to polish the writing. However, none of the paragraphs is written by LLMs directly, and all research ideas are independently proposed by authors without any AI assistance. Claude-Opus-4.1 and Sonnet are used to create figure generation code for Figure 5 and 11. Cursor is included to assist coding, but all generated code is then carefully inspected by authors. Other uses of LLMs in data curation and synthesis are clearly stated in the paper.