

Research Article

From Informal to Formal – Incorporating and Evaluating LLMs on Natural Language Requirements to Verifiable Formal Proofs

Jialun Cao¹, Yaojie Lu²

1. Hong Kong University of Science and Technology, Hong Kong; 2. Institute of Software, Beijing, China

The research in AI-based formal mathematical reasoning has shown an unstoppable growth trend. These studies have excelled in mathematical competitions like IMO and have made significant progress. However, these studies intertwined multiple skills simultaneously—problem-solving, reasoning, and writing formal specifications—making it hard to precisely identify the LLMs’ strengths and weaknesses in each task. This paper focuses on formal verification, an immediate application scenario of formal reasoning, and breaks it down into sub-tasks. We constructed 18k high-quality instruction-response pairs across five mainstream formal specification languages (Coq, Lean4, Dafny, ACSL, and TLA+) in six tasks by distilling gpt-4o and evaluated against ten open-sourced LLMs, including recent popular DeepSeek-R1. We found that LLMs are good at writing proof segments when given either the code, or the detailed description of proof steps. Also, the fine-tuning brought about a nearly threefold improvement at most. Interestingly, we observed that fine-tuning with formal data also enhances mathematics, reasoning, and coding capabilities. Fine-tuned models are released to facilitate subsequent studies at <https://huggingface.co/fm-universe>.

1. Introduction

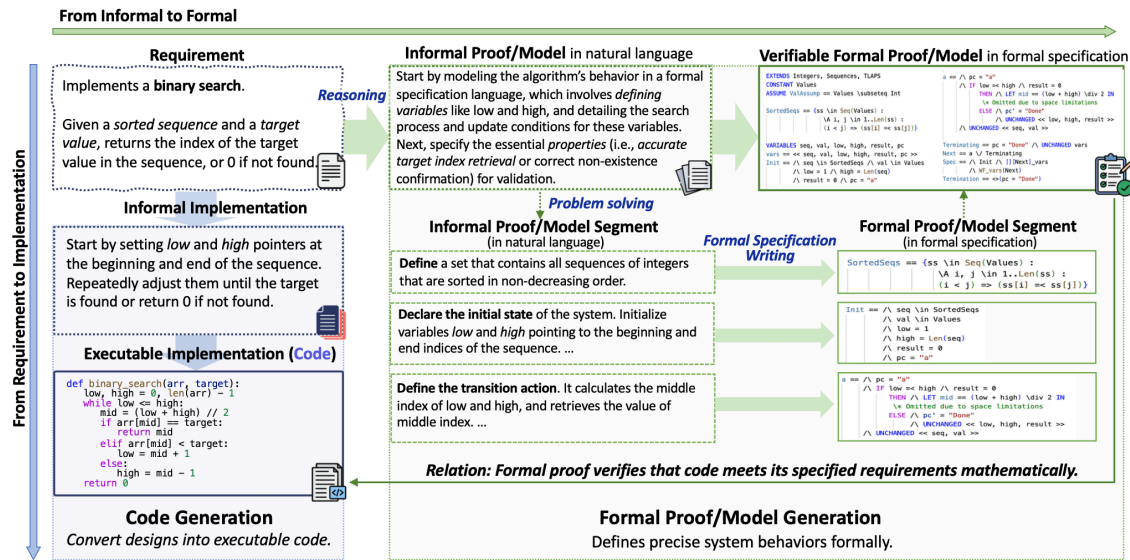


Figure 1. The Illustration of Formal Proof Generation and Its Relation with Code Generation.

As AI-based formal mathematical reasoning reached an inflection point^[41], significant attention and progress in this field have been observed. AlphaProof^[42] achieved silver medal level in the International Mathematical Olympiad (IMO), AlphaGeometry^[43] specialized in proving Euclidean geometry theorems. As reported, the number of publications in this field nearly doubled in 2023, indicating an unstoppable growth trend^[44]. As Fields Medalist Terence Tao imagined, “In the future, instead of typing up our proofs, we would explain them to some GPT”^[51]. However, most current benchmarks cannot precisely reflect the capability to convert informal proofs or requirements in natural language into formal proofs. Most of these benchmarks take mathematical problems^{[2][3][6]} or theorems to be solved^{[7][8][9]} as input, and informal or formal proofs (or parts of proofs) as output. However, these end-to-end benchmarks assess multiple capabilities (e.g., problem-solving, mathematical reasoning, formal

specification writing) in an intertwined manner, making it difficult to isolate and observe LLMs' true capabilities in writing formal proofs or models for verification.

Therefore, we **break down the process** from informal requirements to formal verifiable proof, as shown in Figure 1. Inspired by the code generation (shown in blue) which translates a description of implementation into executable code^{[10][11]}, the formal reasoning process (shown in green) can be seen as translating an **informal requirement** into a **verifiable formal proof** or checkable formal model¹. Particularly, we decompose this process and formulate six tasks (Figure 2). By doing so, the intertwined capabilities can be separated and individually assessed, providing a clearer understanding of LLMs' strengths and weaknesses in each task.

Scope and Targets – We focus on *formal verification*^{[12][13][14][15]} because it is an immediate application scenario of formal mathematical reasoning and the correctness of the output can be verified mechanically. In this paper, we mainly explore four research questions (RQs):

RQ1. How well do LLMs perform in various formal verification tasks? After decomposing the formal verification task into subtasks, we explore LLMs' initial performance in these tasks with zero-shot and few-shot, investigating the strengths and weaknesses that vary between LLMs and tasks.

RQ2. Do LLMs show variability in their capability across different formal specification languages? When mathematicians and proof engineers consider using LLMs to assist in formal verification, they often face uncertainty about which formal specification language is best supported by LLMs. This RQ is designed to provide hints on it.

RQ3. Can fine-tuning improve LLMs' performance in formal verification? Although recent efforts have been made to fine-tune models^{[16][17]}, these LLMs are typically fine-tuned with single formal languages instead of multi-lingual (e.g., combining Coq, Lean, etc.)^[1]. Therefore, we instruction fine-tuned^[17] ^[18] three base LLMs to see whether our constructed fine-tuning dataset FM-ALPACA could improve their capability in formal verification tasks.

RQ4. Can fine-tuning with formal verification data benefit other related tasks (mathematics, reasoning, code)? As recent works have shown LLMs' potential transferability of skills^[19] we thus extend our study to see if models fine-tuned on formal data could show enhanced capabilities in mathematics, reasoning, and coding.

To facilitate the study, we constructed *18k high-quality instruction-response pairs* across five formal specification languages (i.e., Coq, Lean4, Dafny, ACSL, and TLA+) in six formal-verification-related tasks by distilling gpt-4o inspired by prior work^{[46][20][21]}, then split them into 14k instruction fine-tuning data (FM-ALPACA) and 4k benchmarking data (FM-BENCH). In particular, we provide *executable contexts* for all these formal specifications and *automated validation scripts* to validate the correctness of the generated formal proofs inspired from the prior work's artifact preparation^[22]. Finally, we release the fine-tuned LLMs based on three base models at <https://huggingface.co/fm-universe>.

Interestingly, there has been recent discussion on the topic of domain transfer^[1], particularly the transfer of knowledge from other domains such as coding and reasoning to formal domains in order to increase LLMs' reliability^[23], and the anticipated potential of AI in enhancing formal verification processes to support mathematical proofs^{[5][24]}. Our experimental results could potentially provide empirical support for these hypotheses or offer directions for further experimental inquiries.

The contribution of this paper includes:

- **Problem Formulation:** We decompose the formal verification process into six essential tasks. By doing so, the intertwined capabilities can be separated and individually assessed, providing a clearer understanding of LLMs' strengths and weaknesses in each task.
- **Dataset and Benchmark:** We constructed 18k high-quality instruction-response pairs across five mainstream formal specification languages (i.e., Coq, Lean4, Dafny, ACSL, and TLA+) in six formal-verification-related tasks by distilling gpt-4o. They are split into a 14k+ fine-tuning dataset FM-ALPACA and a 4k benchmark FM-BENCH.
- **Executable context and automated validation mechanism:** We provide a Docker container equipped with necessary scripts to facilitate the evaluation of FM-BENCH, significantly lowering the entry barrier for this scenario and making subsequent contributions easier.
- **Insight and Vision:** We fine-tuned several models on FM-ALPACA and observed promising benefits to not only the formal verification tasks, but also mathematics, reasoning, and coding. Our experimental results provide empirical support for the potential of LLMs' capability transfer and hope to shed some light on future research.

2. Task Formulation

Figure 2 illustrates the six sub-tasks. We elaborate on them in detail as follows.

Task 1. Requirement Analysis (abbrev. ReqAna). Requirement analysis^{[25][26][27][28]} is a critical and long-standing research area in software engineering. It facilitates collecting, identifying, categorizing and modeling the users' needs and expectations using various techniques^{[29][30][31][32][33][34]}. In this paper, the requirements are the descriptions in natural language (English)^[33] that details the requirements of the verification/modeling goal and an overall description of the proofs/models. The task is to analyze and break down the final goal into detailed steps described in natural language. The natural language used in this paper is English.

Task 2. Full Proof Generation (abbrev. ProofGen). This task formalizes a requirement in natural language into verifiable proofs or models written in formal specification languages, similar task formulation to existing works^{[35][36][36]}.

Task 3. Proof Segment Generation (abbrev. SegGen). Unlike ProofGen, which requires generating complete proofs/models, SegGen provides more detailed descriptions in natural language and requires LLMs to write less. Given a text description articulating how to implement the proofs/modeling, the task outputs a segment written in the formal specification that serves as a component in the complete proof/model. This task formulation is similar to prior work^{[16][37][38]} and similar to the formulation of code generation^{[39][40][8][6]}.

Task 4. Proof Completion (abbrev. ProofComplete). Similar to code completion^{[41][42][43][44]}, ProofComplete suggests the suffix of the given prefix, similar to prior work^[45]. Note that in order to prevent LLMs from deviating from the original verification goal, we also provide the requirement in our evaluation, although it is not compulsory for this task formulation.

Task 5. Proof InFilling (abbrev. ProofInfill). Given a proof/model with a mask in the middle, the task requires LLMs to fill proper formal specifications so that the completed proofs/models can pass the verifier. This formulation is the same as code infilling^[46]. Also, similar to ProofComplete, we provide the requirement in our evaluation during the infilling to prevent LLMs from deviating from the original verification goal.

Task 6. Proof Generation from Code (abbrev. Code2Proof). In addition to generating formal specifications from natural languages, formal specifications can also be generated from code if the verification goal is the property of a given program. In this paper, we focus mainly on specifications in form of code annotations^{[47][48]}, expressing specifications (e.g., pre-/post-condition, loop invariants) that help one to verify that (part of) a program satisfies certain properties. The task takes the code with properties to be verified as input and outputs the code with generated annotated formal specifications. Similar task formulation can be found in recent works^{[49][50]}.

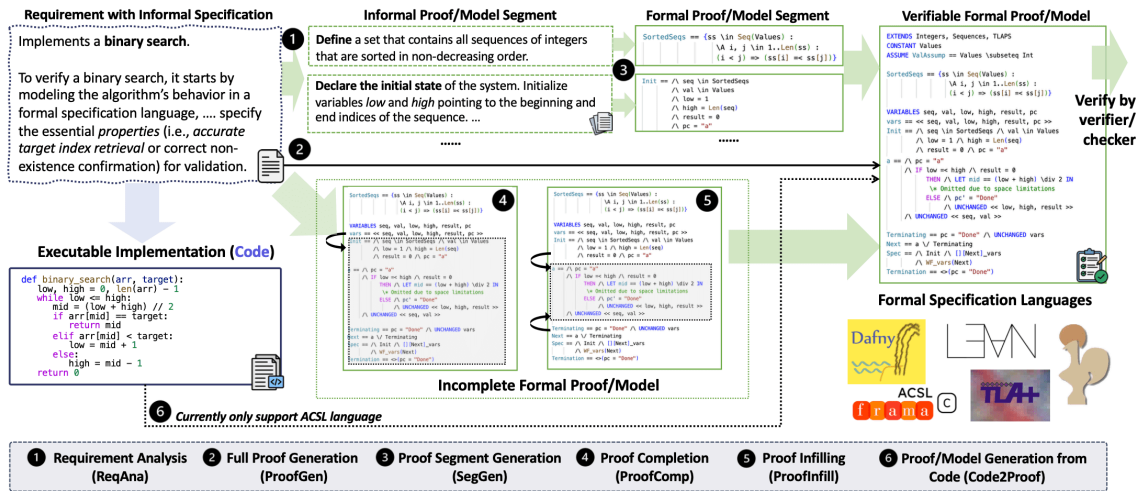


Figure 2. Six tasks towards Informal to Formal Verification

3. Data Construction

3.1. Formal Specification Language Selection

In this study, we consider five formal specification languages that can be used for formal verification, including Coq^[51], Dafny^[52], Lean₄^[53], ACSL (ANSI/ISO C Specification)^[47] and TLA+^{[54][55]}. We selected them in order to cover various *verification paradigms* (i.e., theorem proving and model checking) and *verification scenarios* (e.g., mathematical reasoning and program verification).

First, for *interactive theorem provers* which are suitable for developing rigorous mathematical proofs, we consider Coq^[51] and Lean4^[53] because Coq has been extensively used in academia and research for proving mathematical theorems and in formal verification of software for a long history, while Lean4 garnered considerable attention from the mathematical community^{[16][24][56]} recently. Second, for *programming languages with built-in specification*, we consider Dafny^[52] and ACSL^{[47][57]} because they seamlessly integrate specifications (e.g., pre-/post-conditions, loop invariants) within the code, ensuring the correctness through embedded assertions and conditions. Lastly, for *model checking*^{[58][59]}, we consider TLA+^[54] since it is a representative math-based formal language for modeling algorithms and programs such as concurrent and distributed systems.

3.2. Data Preparation

The workflow of data preparation for FM-ALPACA and FM-BENCH is illustrated in Figure 3. The workflow begins with the data collection, where formal proofs in the desired formal specification languages and related configurations and dependencies are gathered from open-source repositories in Github. Then, formal proofs are **extracted** from the collected repositories. Next, the proofs go through the data quality assurance check by execution, the proofs that cannot be verified successfully are filtered out. The remaining ones are **split** into segments (e.g., definition of functions or conditions).

Given the impracticality of manually writing descriptions for all the collected formal proofs, we leveraged *distilled GPT4*^[60] to generate high-quality informal proof descriptions via meticulous prompting. This alternative is well-established and frequently employed in prior literature^{[21][16][20][61]}. Specifically, for each formal specification language, we designated the model as an expert in that particular language, equipping it with comprehensive domain knowledge about the language’s specifications, essential grammatical cues, and three-shot examples featuring proof segments in the formal language as inputs and natural language descriptions as outputs. This approach ensures that the collected descriptions are of high quality and well-organized. It’s important to note that we did not generate descriptions for proof segments shorter than two lines (such as package imports or constant definitions) because their meaning is self-explained, with the exception of ACSL, whose proof segments are typically 1-2 lines.

After the descriptions for both full and segment proofs were prepared, we then prepared the data pairs for each task as shown in **Task-wise Data Pairing** in Figure 3. Note that for Task 4, i.e., Proof Completion, to prepare the incomplete formal proof, we randomly choose a line number and delete the lines in the proof after the line. For Task 5, i.e., Proof Infill, we randomly choose two line numbers and mask the lines between them. In case the remaining lines of proof cannot provide sufficient information for the proof generation, we also provide the informal proof for these two tasks.

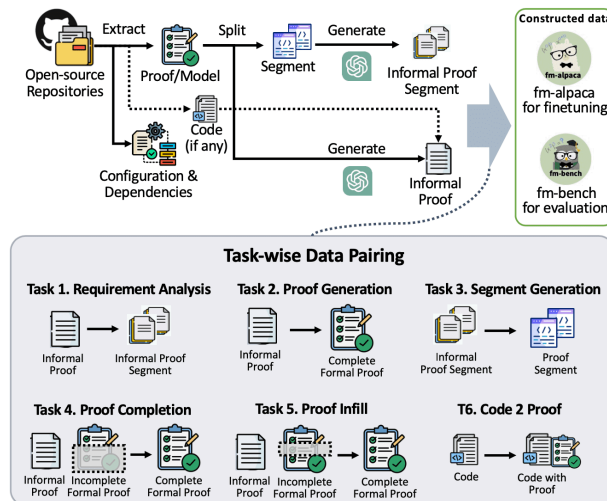


Figure 3. The Illustration of Data Preparation.

After pairing the instruction–response for different tasks, we manually designed five task instructions for each task and randomly assigned one for each paired data to increase instruction diversity and avoid overfitting to certain instructions^{[62][63][64]}.

3.3. Data Statistics

The specification–language-wise and task-wise statistics are shown in Table 1 and Table 2. In particular, Table 1 presents a detailed breakdown of the number of proofs and segments across five specification languages. Note that we split the prepared data for all the tasks and specification languages into

an 8:2 ratio, *i.e.*, 80% for fine-tuning, named FM-ALPACA, 20% for benchmarking, named FM-BENCH, and show the separate statistics. In particular, there are 4k+ verifiable proofs in total, with 249 ~ 2k proofs for each language. These proofs were split into 18k+ proof segments, with an average of 3.6k segments for each language. The reason why the ratio of segments in FM-ALPACA and FM-BENCH is slightly less than 8:2 is that the train-test split was applied to proofs, while the number of split segments in each proof varies.

	Num of Proofs			Num of Segments		
	Total	FM-ALPACA	FM-BENCH	Total	FM-ALPACA	FM-BENCH
Coq	2126	1683	443	14939	11638	3301
Lean4	1163	919	244	1578	1261	317
ACSL	544	426	118	765	598	167
Dafny	249	206	43	417	348	69
TLA+	256	199	57	594	476	118
Total	4338	3433	905	18293	14321	3972
Average	868	687	181	3659	2864	794

Table 1. Formal-Specification-Language-wise Statistics of Formal Verification Data

Table 2 shows the task-wise statistics. There are 18k instructions across six tasks, with an average of 3k instructions for each task. After splitting the train-test set, FM-ALPACA contains 14k, and FM-BENCH has nearly 4k instructions. It is clear that the number of instructions for the task Segment Proof Generation (SegGen) is far more than that for Requirement Analysis (ReqAna) and Full Proof Generation (ProofGen) because one full proof can be split into numerous pieces of proof segments, and one proof can contribute to only one instruction for ReqAna and ProofGen. Note that the number of ReqAna (627) and ProofGen (700) is unequal because we filtered out the instructions with more than 2048 tokens considering the context limits.

	Task	Total	FM-ALPACA	FM-BENCH
1	Requirement Analysis	627	496	131
2	Full Proof Generation	700	557	143
3	Segment Proof Generation	14843	11597	3246
4	Proof Complete	658	520	138
5	Proof Infill	1439	1146	293
6	Code2Proof	70	56	14
	Total	18337	14372	3965
	Average	3056	2395	661

Table 2. Task-wise Statistics of Formal Verification Data

3.4. Validation Mechanism

For the tasks whose outputs are written in formal specification languages, we verify the full proofs against the corresponding verifiers, *i.e.*, Coq, Dafny, Lean4 use their own proving environment; formal specification written in TLA+ can be checked by TLC^[65]; C programs with ACSL specifications can be checked by Frama-C^{[57][66]}. Also, for the proof segments that cannot be verified independently, for each extracted segment, we prepared a proof template with a placeholder during data preparation. Whenever a generated segment is to be verified, we replace the placeholder in the template with the segment and verify the completed formal proof.

For the task whose outputs are written in natural language (i.e., ReqAna), we calculate the Bleu score^[67] between the descriptions in FM-BENCH with the predicted outputs.

LLMs	Size	Proof Generation				Proof Segment Generation				ProofComplete				ProofInfill				Cd2Prf	Ave			
		TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	TLA	Coq			Lean	Dafny	ACSL
W/o Fine-tuning																						
llama3.1-instruct	8B	0.00	4.26	1.69	0.00	14.29	1.43	1.19	6.06	8.33	0.00	0.00	3.57	5.56	28.57	0.00	4.85	7.69	13.33	50.00	71.43	11.11
llama3.1-instruct	70B	7.69	6.38	10.17	20.00	7.14	22.86	4.86	13.64	11.11	39.68	27.27	3.57	13.89	28.57	0.00	7.27	24.62	13.33	21.43	21.43	15.25
qwen2.5-instruct	7B	0.00	2.13	5.08	0.00	21.43	1.43	1.05	7.58	5.56	1.59	0.00	13.10	19.44	28.57	0.00	5.45	18.46	13.33	50.00	57.14	12.57
qwen2.5-instruct	72B	7.69	6.38	18.64	20.00	21.43	12.86	3.29	12.12	25.00	12.70	0.00	11.90	33.33	28.57	0.00	9.70	24.62	26.67	35.71	57.14	18.39
qwen2.5-coder-instruct	7B	0.00	2.13	6.78	0.00	0.00	2.86	2.14	12.12	11.11	3.17	0.00	4.76	11.11	14.29	0.00	3.64	23.08	20.00	50.00	57.14	11.22
qwen2.5-coder-instruct	32B	0.00	6.38	11.86	0.00	21.43	12.86	4.07	9.85	25.00	26.98	27.27	5.95	36.11	42.86	2.94	7.88	33.85	20.00	71.43	57.14	21.19
starcoder-instruct	15B	7.69	2.13	8.47	30.00	28.57	27.14	1.97	6.06	27.78	0.00	27.27	14.29	30.56	42.86	0.00	11.52	27.69	20.00	35.71	28.57	18.27
deepseek-coder-instruct-v1.5	7B	0.00	2.13	6.78	0.00	7.14	4.29	1.73	3.79	5.56	0.00	9.09	2.38	13.89	28.57	0.00	3.03	13.85	33.33	7.14	14.29	7.85
deepseek-coder-instruct	33B	0.00	0.00	1.69	0.00	0.00	2.86	1.94	4.55	11.11	3.17	0.00	3.57	13.89	42.86	0.00	2.42	23.08	6.67	21.43	7.14	7.32
deepseek-R1	671B	15.38	8.51	25.42	30	35.71	22.86	10.7	21.21	22.22	49.21	45.45	17.86	36.11	28.57	5.88	13.33	40	6.67	42.86	64.29	27.11
Task-wise Average		8.65				10.61				18.44				17.38				43.57				
W/ Fine-tuning																						
llama3.1-fma	8B	0.00	6.38	8.47	20.00	57.14	41.43	29.98	21.97	25.00	90.48	36.36	8.33	11.11	28.57	0.00	0.00	7.69	0.00	21.43	21.43	21.79
llama3.1-ultrachat	8B	0.00	0.00	1.69	0.00	7.14	0.00	0.00	3.79	0.00	0.00	0.00	5.95	13.89	14.29	0.00	4.85	9.23	0.00	14.29	0.00	3.76
llama3.1-ultrachat-fma	8B	0.00	8.51	10.17	30.00	50.00	41.43	35.72	29.55	33.33	95.24	18.18	8.33	11.11	28.57	2.94	0.00	7.69	0.00	28.57	35.71	23.75
llama3.1-tulu	8B	0.00	0.00	1.69	0.00	7.14	0.00	0.71	6.06	0.00	0.00	0.00	5.95	8.33	28.57	0.00	1.82	10.77	0.00	57.14	50.00	8.91
llama3.1-tulu-fma	8B	0.00	4.26	11.86	30.00	35.71	42.86	36.71	27.27	36.11	98.41	18.18	10.71	11.11	42.86	2.94	2.42	6.15	0.00	35.71	50.00	25.16
qwen2.5-fma	7B	0.00	4.26	11.86	10.00	71.43	38.57	27.95	27.27	22.22	87.30	27.27	9.52	13.89	28.57	0.00	0.00	7.69	0.00	35.71	21.43	22.25
qwen2.5-coder-fma	7B	0.00	6.38	18.64	20.00	35.71	44.29	36.50	34.09	33.33	98.41	36.36	9.52	16.67	28.57	5.88	1.82	13.85	13.33	42.86	50.00	27.31
deepseek-coder-v1.5-fma	7B	0.00	2.13	16.95	0.00	21.43	34.29	25.30	31.06	25.00	84.13	36.36	8.33	19.44	28.57	0.00	1.82	7.69	6.67	21.43	35.71	20.32
Task-wise Average		12.72 ↑				33.64 ↑				17.92				9.31				22.92				

Table 3. RQ1-3: Pass@1 Accuracy of LLMs' Performance Across Formal Verification Task and Formal Specification Languages with (w/) and without (w/o) fine-tuning. The greener, the better.

- * -fma: fine-tuned with FM-ALPACA.
- * -ultrachat: fine-tuned with UltraChat.
- * -tulu: fine-tuned with Tulu3.
- * -ultrachat-fma: fine-tuned with both UltraChat and FM-ALPACA.
- * -tulu-fma: fine-tuned with both Tulu3 and FM-ALPACA.

LLMs	Size	TLA	Coq	Lean	Dafny	ACSL
w/o fine-tune						
llama3.1-instruct	8B	0.33	0.32	0.4	0.47	0.34
llama3.1-instruct	70B	0.38	0.31	0.43	0.5	0.37
qwen2.5-instruct	7B	0.34	0.29	0.35	0.41	0.26
qwen2.5-instruct	72B	0.35	0.31	0.36	0.36	0.24
qwen2.5-coder-instruct	7B	0.37	0.34	0.42	0.44	0.31
qwen2.5-coder-instruct	32B	0.34	0.31	0.36	0.41	0.26
starcoder-instruct	15B	0.39	0.36	0.46	0.33	0.37
deepseek-coder-instruct-v1.5	7B	0.42	0.37	0.47	0.55	0.43
deepseek-coder-instruct	33B	0.43	0.36	0.5	0.54	0.46
deepseek-R1	671B	0.32	0.27	0.33	0.27	0.28
Average		0.38				
w/ fine-tune						
llama3.1-fma	8B	0.53	0.28	0.58	0.67	0.5
llama3.1-ultrachat	8B	0.44	0.34	0.47	0.57	0.5
llama3.1-ultrachat-fma	8B	0.55	0.42	0.62	0.65	0.68
llama3.1-tulu	8B	0.44	0.38	0.49	0.52	0.47
llama3.1-tulu-fma	8B	0.64	0.41	0.64	0.7	0.76
qwen2.5-fma	7B	0.57	0.24	0.61	0.63	0.57
qwen2.5-coder-fma	7B	0.59	0.31	0.62	0.75	0.75
deepseek-coder-v1.5-fma	7B	0.51	0.36	0.55	0.64	0.46
Average		0.54 (41% ↑)				

Table 4. Evaluation on Requirement Analysis

- * -fma: fine-tuned with FM-ALPACA.

4. Experiments

4.1. Experiment Setup

Studied LLMs. We selected ten LLMs as baselines without fine-tuning, including llama3.1-instruct-8B/70B^[68], qwen2.5-instruct-7B/72B^[69], qwen2.5-coder-instruct-7B/-32B^[70], starcoder-instruct-15B^[71], deepseek-coder-instruct-7B-v1.5, deepseek-coder-instruct-33B^[72], and deepseek-R1^[73]. Note that we avoid evaluating the GPT-series LLMs by OpenAI because the descriptions in FM-BENCH were generated by GPT-4o, making the evaluation fairer.

Fine-tuning. Instruction fine-tuning^{[17][18][20][74]} aims to improve a model’s ability to effectively respond to human instructions and has shown strong experimental potential in model enhancement. We select llama3.1-8B^[68], qwen2.5-7B^[69], and deepseek-coder-7B-v1.5^[72] as base models for fine-tuning. We selected these three models because they have shown promising capability in tasks such as coding, mathematics, and reasoning, and fine-tuning models in their scale is relatively affordable compared with fine-tuning larger scale models. We fine-tuned the three aforementioned base models over three epochs using a learning rate $2e-5$, a warm-up ratio of 0.04, a batch size of 512, and a cosine learning rate scheduler.

Baseline Fine-tuning Datasets: To distinguish whether the capability improvement is simply because more instruction tuning is applied, we also include two commonly used fine-tuning datasets for comparison. We select UltraChat^[20] and Tulu-V3^[75] as baseline fine-tuning datasets for their popularity. In particular, UltraChat is a large-scale dataset of instructional conversations that contains 1.5 million high-quality multi-turn dialogues and covers a wide range of topics and instructions. Tulu-v3^[75] embraces new data that is either carefully manually curated for quality or generated from GPT models. It is an enhancement of its previous versions^{[74][61]}, focusing more on core skills of knowledge recall, reasoning, mathematics, coding, instruction following, general chat, and safety.

Benchmarks for Related Capabilities (RQ4). To comprehensively evaluate the model’s capabilities, we tested the fine-tuned models on a series of benchmarks: Math^[76] and GSM-8K^[77] for mathematical reasoning, BBH^[78] for general reasoning, HumanEval^[40] and MBPP^[79] for coding.

Inference Strategies. We adopt different settings for different RQs. In particular, We use (1) the greedy sampling strategy to generate one single greedy sample with a temperature of 0.0 and calculate Pass@1, and (2) nucleus sampling^[80], where five solution samples were randomly generated with a temperature of 0.2 for RQ1 and RQ2. We also consider different in-context learning strategies, including zero-shot and few-shot (we used 3-shot in the experiment). For RQ3, we use a zero-shot greedy search with a temperature of 0.0 and a few-shot nuclear search with a temperature of 0.2 for a fair comparison.

Experiment Environment. The fine-tuning experiment was conducted on 32 Nvidia A100-40G GPUs, while inference was on a single Nvidia A100-80G GPU with vLLM^[81].

4.2. RQ1. Basic Performance across Formal Specification Tasks

To understand the current LLMs’ performance in six tasks, we evaluate 8 LLMs against FM-BENCH with model size ranges from 7B to 72B. The upper part of Table 3 and the upper part of Table 4 show LLMs’ basic performance without fine-tuning.

Task-wise: LLMs perform the best in generating proof from code (Code2Proof), with an average of 43.57% Pass@1, followed by ProofComplete (18.44%) and ProofInfill (17.38%). In contrast, LLMs fall short in generating both the entire formal proof (8.65%) and the proof segments (10.61%). We analyzed the failures and found that syntax errors account for a large proportion, with 12.15% failures caused by syntax errors (Appendix B). The observation echoes the motivation of prior work^[16] and is reasonable due to the grammar difference between most formal specification languages and other programming languages like Python. Regarding requirement analysis, as shown in the upper part of Table 4, the Bleu scores between the ground-truth description and LLM-generated ones range from 0.24 to 0.55.

LLM-wise: Without fine-tuning, DeepSeek-R1 achieved the best average (27.11%), followed by qwen2.5-coder-instruct-32B (21.19%).

Model Size: Larger LLMs generally perform better than smaller LLMs. For example, llama3.1-8B only achieved 1.43% in generating TLA+ segments, while llama3.1-70B boosts to 22.86% in the same task. However, there are several exceptions worth noticing, especially for ProofInfill and Code2Proof. For example, llama-3.1-8B achieved 50% in ProofInfill (ACSL), yet the performance drops to 21.43% using the 70B model. Similar observations can be found in qwen2.5-instruct. The decrease in performance is inherently due to the fine-tuning strategy of these instruction models: they are trained to excel in generating rather than filling in the blanks^[46]. Also, we conducted a more detailed examination of generated segments and observed that larger LLMs tend to fill in the proof segments that *not only look more plausibly correct and well-organized but also include extra content*. The additional content, yet, is either redundant, as it repeats information that appears in the subsequent proof, or is incomplete. Promisingly, recent model developers have noticed such conundrums and refined their fine-tuning strategy for fill-in-the-middle tasks^[72].

LLMs	Size	TLA			Coq			Lean			Dafny			ACSL		
		Zero-shot		Few-Shot	Zero-shot		Few-Shot	Zero-shot		Few-Shot	Zero-shot		Few-Shot	Zero-shot		Few-Shot
		Pass@1	Pass@5	Pass@1	Pass@1	Pass@5	Pass@1	Pass@5	Pass@1	Pass@5	Pass@1	Pass@5	Pass@1	Pass@1	Pass@5	Pass@1
llama3.1-instruct	8B	1.43	4.29	11.43	1.19	1.90	10.12	6.06	8.33	9.85	8.33	11.11	19.44	0.00	3.17	38.10
llama3.1-instruct	70B	22.86	27.14	38.57	4.86	6.01	3.46	13.64	17.42	15.91	11.11	13.89	16.67	39.68	52.38	76.19
qwen2.5-instruct	7B	1.43	2.86	5.71	1.05	1.39	12.80	7.58	9.09	9.85	5.56	5.56	13.89	1.59	3.17	66.67
qwen2.5-instruct	72B	12.86	17.14	24.29	3.29	3.80	21.05	12.12	14.39	19.70	25.00	25.00	25.00	12.70	15.87	90.48
qwen2.5-coder-instruct	7B	2.86	5.71	12.86	2.14	2.99	14.30	12.12	13.64	18.94	11.11	11.11	19.44	3.17	4.76	87.30
qwen2.5-coder-instruct	32B	12.86	15.71	30.00	4.07	4.96	18.40	9.85	15.15	13.64	25.00	27.78	30.56	26.98	33.33	96.83
deepseek-coder-instruct-v1.5	7B	4.29	8.57	4.29	1.73	2.82	8.79	3.79	5.30	17.42	5.56	5.56	19.44	0.00	4.76	60.32
deepseek-coder-instruct	33B	2.86	2.86	18.57	1.94	3.23	11.27	4.55	6.82	22.73	11.11	13.89	22.22	3.17	25.40	92.06
Language-wise Average		12.14			6.15			12.00			15.97			34.92		
Average		7.68	10.54	18.22	2.53	3.39	12.52	8.71	11.27	16.01	12.85	14.24	20.83	10.01	17.86	75.99
Improvement Ratio			37.15%↑	99.98%↑		33.70%↑	360.58%↑		29.31%↑	54.37%↑		10.82%↑	51.33%↑		63.64%↑	532.83%↑

Table 5. RQ2: Language-wise LLMs' Performance. Pass@1 and Pass@5 Accuracy in Generating Proof Segments Across Formal Specification Languages Under Zero/Few-shot without fine-tuning.

*The improvement ratios shown in "Pass@5" column are calculated by comparing with the scores in Pass@1, and the ratios shown in "Few-shot" column are calculated by comparing with the scores in Pass@5.

Fine-tuning Dataset	MATH			Reasoning		Coding			Average							
	MATH	GSM-8K	Average	bbh	HumanEval	MBPP	Average									
ultrachat	17.54	61.33	39.44	62.64	19.51	36.4	27.96	39.48								
UltraChat + fma	16.16	7.87%↓	62.32	1.61%↑	39.24	0.49%↓	62.14	0.80%↓	19.51	62.53%↑	35.2	3.30%↓	33.46	19.67%↑	41.51	5.14%↑
tulu3	27.36	75.82	51.59	62.47	64.02	48.8	56.41	55.01								
tulu3 + fma	29.48	7.75%↑	75.44	0.50%↓	52.46	1.69%↑	63.16	1.10%↑	64.63	0.95%↑	49.4	1.23%↑	57.02	1.07%↑	55.76	1.37%↑

Table 6. RQ4: Capability Migration from FM-ALPACA to Math, Reasoning, and Coding.

4.3. RQ2. Formal Specification Languages-wise Capability

Table 5 shows the LLMs' performance across formal specification languages in the task of generating proof segments (SegGen). This task accounts for the most instructions and serves as the basic capability for other proof generation tasks. We can see that LLMs perform the best in ACSL (average: 34.92%), followed by Dafny (15.97%) while performing unsatisfactorily in other formal specification languages. The observation is reasonable because the syntax of ACSL is basically an annotation of C language, while Dafny shares similar grammar as C# and Java. Thus, generating proof segments in ACSL and Dafny is generally easier than generating other specification languages.

In addition, we explore whether increasing the attempts (1 → 5) with a higher temperature (0.0 → 0.2) and in-context learning could bring about improvement. The improvement ratios are shown in red in Table 5. The results of Pass@5 are better than those of Pass@1, with an average score increase from 10.82% (Dafny) to 63.64% (ACSL) in different languages. Moreover, when using 3-shot, the performance increases dramatically, with 51.33% (Dafny) to over five times (ACSL) improvement compared with zero-shot Pass@5. The results indicate the potential of in-context learning in generating correct specification languages.

4.4. RQ3. Improvement by Fine-tuning

We further investigate whether FM-ALPACA could bring about improvement. The lower part of Table 3 and Table 4 shows the results. From Table 3, dramatic improvements can be observed in generating full and segmental proofs after fine-tuning. Note that the model size of fine-tuned models is 7B ~ 8B, while the performance largely outperforms the 70B+ models without fine-tuning. Furthermore, after fine-tuning with formal data, the 7 ~ 8B fine-tuned models can achieve comparable or slightly better performance than Deepseek-R1-671B, with 27.31% achieved by qwen2.5-coder-7B fine-tuned with FM-ALPACA (R1-671B: 27.11%). It may suggest the possibility of distilling domain-specific small models for handier usage.

Task-wise: Improvements in generation tasks (i.e., ProofGen, SegGen, and ProofComplete) are substantial. ProofGen doubles the performance, and SegGen more than triples. The dramatic increases happen in all models fine-tuned with FM-ALPACA in SegGen Task, from nearly all zeros to 29.98% ~ 90.48%. An increase of 41% can also be observed in Table 4. The experimental improvements make evident the effectiveness of fine-tuning in formal verification tasks.

Yet, drops can be observed in fill-in-the-middle tasks (i.e., ProofInfill and Code2Proof). The results echo the observation made in RQ1 (Section 4.2), where the large LLMs perform worse than small LLMs in fill-in-the-middle tasks. The results also indicate the necessity of adopting different fine-tuning strategies other than instruction tuning only.

Fine-tuning Datasets: Take a closer look at the LLMs fine-tuned with general-purpose datasets (*i.e.*, llama3.1-ultrachat and llama3.1-tulu) in Table 3, with them only, no or opposite effects can be observed. The results indicate the *complementarity* of FM-ALPACA and existing general-purpose fine-tuning datasets. Additionally, by combining with general-purpose datasets, the performance can be *further improved* (*e.g.*, llama3.1-tulu-fma).

Comparison with Few-shot: Compared with the best results in Table 5 achieved by 3-shot, the results after fine-tuning (Table 3) still generally outperform the 3-shot results. The results indicate that although in-context learning can improve LLMs' performance, the enhancement is limited. *Further significant improvements still require fine-tuning with formal data.* This may also suggest that in-context learning alone cannot adequately address capability deficits in formal verification tasks but rather stem from a lack of knowledge.

4.5. RQ4. Capability Migration from Formal Verification to Related Tasks

Finally, we explore whether fine-tuning with FM-ALPACA could benefit related capabilities. Table 6 shows the results. The base model is llama3.1-8B, fine-tuned under two base fine-tuned datasets with and without FM-ALPACA. On average, with FM-ALPACA, an increase of 1.37% to 5.15% can be observed. Interestingly, a dramatic increase (62.53%) can be observed in HumanEval compared with the performance of the model that is only fine-tuned with UltraChat. The experiment may indicate that feeding more formal data may improve LLMs' coding, reasoning, and math capabilities.

5. Conclusion

This paper contributes a comprehensive assessment and formulation to understand LLMs' capability in formal verification. We constructed 18k high-quality instruction-response pairs across five formal specification languages in six tasks. The fine-tuned models, fine-tuning data, and the benchmark are released to facilitate subsequent studies.

Limitations

This paper has two primary limitations that offer avenues for future research. First, the primary limitation of our work is that our benchmark relies on model-generated data. While this approach effectively reduces manual efforts; it may introduce biases and data leakage issues in the dataset towards the models that generated the data. To address this limitation, we use gpt-4o to generate the natural language descriptions, while during the evaluation, we use other LLMs for evaluation. Second, another limitation of our work lies in the validation design. When creating ProofInfill and ProofComplete data, it is possible that the properties to be verified or theorems to be proven are masked. If LLMs happened not to generate these properties/theorems, the generated "proofs/models" could escape the verifier/checker, mistakenly labeling the output as correct. To avoid this scenario, we include the requirement descriptions as part of the input, guiding LLMs to generate the necessary properties or theorems without omission.

Appendix A. Related Work

The formal specification datasets or benchmarks offer a standard, well-defined set of problems, providing a shared challenge that helps build a community of practice among researchers. According to different verification techniques, the existing benchmarks mainly fall into two categories; we discuss them separately.

A.1. Theorem Proving Datasets

Formal theorem proving represents theorems and proofs in a machine-verifiable format^[82], ensuring their correctness using rigorous logical rules. A recent survey^[44] summarized the existing datasets for theorem proving. In particular, the *informal benchmarks* craft the proofs from various sources such as ProofWiki, textbooks, and public corpus. NL-PS^[83] first builds a natural language premise selection dataset source from ProofWiki. Similarly, NaturalProofs^[61] further incorporates data from Stacks and textbooks, resulting in a dataset with roughly 25k examples. Adapted from it, NaturalProofs-Gen^[81] contains around 14.5k theorems for informal proof generation. Moreover, MATcH^[84] constructs over 180k statement-proof pairs for matching using the MREC corpus².

For *formal datasets*, a line of efforts focuses on extracting and cleaning theorems and proofs written in various specification languages (*e.g.*, Coq, Isabelle, Lean) from established formal libraries and verification projects. For example, LeanDojo^[7] extracts over 98k theorems and proofs with 130k premises from Lean mathlib^[85]. Besides extracting data from existing projects, several works *manually* annotate or formalize the problems in natural language. For example, MiniF2F^[86] manually formalizes 488 Olympiad-level problems across 4 proof systems and equally splits them into a validation set and a test set. FIMO^[87] and ProofNet^[88] formalize the theorem statements of the International Mathematical Olympiad and undergraduate-level problems in Lean. In

addition, datasets for Dafny also attract research contributions because industries like Amazon adopted Dafny to verify cryptographic libraries, authorization protocols, a random number generator, and the Ethereum virtual machine. Dafny datasets such as CloverBench^[40] and DafnyGym^[89].

A.2. Model checking datasets

Model checking is an automated technique used in computer science and formal methods to verify the correctness of systems, particularly those with finite state spaces. It systematically checks whether a system’s model satisfies a given specification, usually expressed in formal specification languages. The basic idea is to explore all possible system states to ensure the desired properties hold in every conceivable scenario.

Model checking benchmarks are less than that for theorem proving. Currently, there are few model-checking benchmarks for proving, while several model-checking subjects are going with specific model-checking languages such as CMurphi^[90] and TLA+^[54]. In particular, CMurphi is a software tool used to verify concurrent and distributed systems through explicit state enumeration. It implements the Murphi verification language, which allows users to describe finite-state systems in a procedural style. The core principle behind CMurphi is to explore the state space of a system exhaustively to check for violations of specified invariants or properties. Another example is TLA+ (Temporal Logic of Actions), a high-level language for modeling programs and systems suitable for concurrent and distributed systems.

Appendix B. Proportion of Failures Caused by Syntax Error

We listed the proportions of failures caused by syntax errors for each LLM and each task in Table 8. We used a set of pre-defined keywords (summarized in Table 7 to identify if a verification failure is caused by syntax errors. Specifically, we consider a failure caused by syntax error if its error message contains at least one keyword in Table 7.

	Language	Keywords
1	Coq	“Syntax Error:”
2	Lean4	“unexpected token”, “unknown identifier”, “type mismatch”
3	ACSL	“unexpected token”
4	Dafny	“type errors detected”
5	TLA+	“***Parse Error***”, “Unknown operator”

Table 7. Keywords for identifying syntax error raised by each language’s verifier.

LLMs	Size	ProofGen					SegGen					ProofComplete				ProofInfill					Code2P
		TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	TLA	Coq	Lean	Dafny	ACSL	
llama3.1-instruct	8B	0/13	1/45	33/58	0/10	1/12	43/69	746/2910	78/124	4/33	4/63	4/11	45/81	16/34	1/5	16/34	44/157	27/60	0/13	1/7	0/4
llama3.1-instruct	70B	2/12	0/44	38/53	1/8	1/13	43/54	71/2802	75/114	16/32	5/38	6/8	44/81	10/31	1/5	15/34	39/153	22/49	0/13	2/11	3/11
qwen2.5-instruct	7B	10/13	1/46	30/56	1/10	0/11	26/69	303/2914	53/122	5/34	6/62	9/11	35/73	8/29	1/5	14/34	66/156	31/53	0/13	4/7	3/6
qwen2.5-instruct	72B	7/12	0/44	25/48	1/8	9/11	47/61	40/2848	54/116	11/27	0/55	7/11	35/74	8/24	1/5	13/34	59/149	20/49	0/11	0/9	1/6
qwen2.5-coder-instruct	7B	9/13	0/46	32/55	0/10	0/14	60/68	22/2882	72/116	6/32	53/61	10/11	12/80	8/32	2/6	17/34	34/159	26/50	1/12	0/7	0/6
qwen2.5-coder-instruct	32B	7/13	0/44	26/52	1/10	2/11	50/61	31/2825	68/119	9/27	25/46	5/8	24/79	6/23	3/4	17/33	51/152	16/43	0/12	0/4	0/6
deepseek-coder-instruct	7B	10/13	0/46	29/55	0/10	0/13	35/67	9/2894	62/127	3/34	43/63	9/10	14/82	11/31	2/5	22/34	18/160	23/56	0/10	1/13	1/12
deepseek-coder-instruct	33B	12/13	0/47	43/58	4/10	0/14	66/68	52/2888	81/126	8/32	45/61	10/11	3/81	6/31	2/4	19/34	6/161	23/50	0/14	0/11	1/13
starcoder-instruct	15B	8/12	0/46	24/54	2/7	3/10	36/51	10/2887	67/124	12/26	1/63	6/8	4/74	7/27	1/4	20/34	6/149	23/49	0/12	0/9	1/10

Table 8. Proportion of Verification Failures Caused by Syntax Errors

¹ Denominator represents the total number of failures.

Appendix C. Example Specifications in FM-ALPACA and FM-BENCH

The examples of the five formal specification languages are shown in Figure 4.

<pre> Proof. intros EM R H. destruct (EM (R true false)). - exists (fun _ => true). - intros []; firstorder. - exists (fun b => b). intros b. split. + reflexivity. + destruct b, y; intros HR; easy now symmetry in HR. Qed. </pre>	<pre> theorem gcd_dvd_left {a b : Int} : (gcd a b : Int) a := by have := Nat.gcd_dvd_left a.natAbs b.natAbs rw [- Int.ofNat_dvd] at this exact Int.dvd_trans this natAbs_dvd_self theorem gcd_dvd_right {a b : Int} : (gcd a b : Int) b := by have := Nat.gcd_dvd_right a.natAbs b.natAbs rw [- Int.ofNat_dvd] at this exact Int.dvd_trans this natAbs_dvd_self </pre>	<pre> #include <stdio.h> /*@ requires a >= 0 && b >= 0; ensures \result == \old(a)*\old(b); assigns \nothing; */ int mul(int a, int b) { int x = a, y = b, prod = 0; /*@ loop invariant prod == (a-x)*y; loop assigns prod, x; */ while(x >= 0) { prod = prod + y; x--; } return prod; } int main() { int pdt = mul(2, 5); //@ assert pdt == 10; } </pre>
<pre> method factorial(n: nat) returns (res: nat) ensures res == fact(n) { var i := 1; res := 1; while i < n + 1 invariant 0 < i <= n + 1 invariant res == fact(i - 1) { res := i * res; i := i + 1; } } </pre>	<pre> PickDifferentColor == /\ BeanCount > 1 /\ can.black >= 1 /\ can.white >= 1 /\ can' = [can EXCEPT !.black = @ - 1] TerminationHypothesis == IF can.white % 2 = 0 THEN <<(can.black = 1 /\ can.white = 0) ELSE <<(can.black = 0 /\ can.white = 1) </pre>	

Figure 4. Formal Specification Languages in FM-bench

Appendix D. Collected Repositories

We listed the repositories that were collected for data construction in the following. Note that one can easily add more repositories into FM-ALPACA and FM-BENCH.

For ACSL:

- <https://github.com/manavpatnaik/frama-c-problems>
- <https://github.com/fraunhoferfokus/acsl-by-example>

For TLA+:

- <https://github.com/tlaplus/Examples>

For Lean4:

- <https://github.com/leanprover/lean4>

For Coq:

- <https://github.com/coq/coq>

For Dafny:

- <https://github.com/vladstjeroiu/Dafny-programs>

Appendix E. Complete Evaluation Result

The Pass@1 and Pass@5 are shown in Table 9. It is a completed version of Table 3.

Pass@1																					
LLMs	Size	ProofGen						SegGen				ProofComplete				ProofInfill			Code2Sep		
		TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	TLA	Coq	Lean	Dafny	ACSL	ACSL
W/o Fine-tuning																					
llama3.1-instruct	8B	0.00	4.26	1.69	0.00	14.29	1.43	1.19	6.06	8.33	0.00	0.00	3.57	5.56	28.57	0.00	4.85	7.69	13.33	50.00	71.43
llama3.1-instruct	70B	7.69	6.38	10.17	20.00	7.14	22.86	4.86	13.64	11.11	39.68	27.27	3.57	13.89	28.57	0.00	7.27	24.62	13.33	21.43	21.43
qwen2.5-instruct	7B	0.00	2.13	5.08	0.00	21.43	1.43	1.05	7.58	5.56	1.59	0.00	13.10	19.44	28.57	0.00	5.45	18.46	13.33	50.00	57.14
qwen2.5-instruct	72B	7.69	6.38	18.64	20.00	21.43	12.86	3.29	12.12	25.00	12.70	0.00	11.90	33.33	28.57	0.00	9.70	24.62	26.67	35.71	57.14
qwen2.5-coder-instruct	7B	0.00	2.13	6.78	0.00	0.00	2.86	2.14	12.12	11.11	3.17	0.00	4.76	11.11	14.29	0.00	3.64	23.08	20.00	50.00	57.14
qwen2.5-coder-instruct	32B	0.00	6.38	11.86	0.00	21.43	12.86	4.07	9.85	25.00	26.98	27.27	5.95	36.11	42.86	2.94	7.88	33.85	20.00	71.43	57.14
deepseek-coder-instruct	7B	0.00	2.13	6.78	0.00	7.14	4.29	1.73	3.79	5.56	0.00	9.09	2.38	13.89	28.57	0.00	3.03	13.85	33.33	7.14	14.29
deepseek-coder-instruct	33B	0.00	0.00	1.69	0.00	0.00	2.86	1.94	4.55	11.11	3.17	0.00	3.57	13.89	42.86	0.00	2.42	23.08	6.67	21.43	7.14
star-coder-instruct	15B	7.69	2.13	8.47	30.00	28.57	27.14	1.97	6.06	27.78	0.00	27.27	11.90	25.00	42.86	0.00	9.70	24.62	20.00	35.71	28.57
W/ Fine-tuning																					
llama3.1-fma	8B	0.00	6.38	8.47	20.00	57.14	41.43	29.98	21.97	25.00	90.48	36.36	8.33	11.11	28.57	0.00	0.00	7.69	0.00	21.43	21.43
llama3.1-ultrachat	8B	0.00	1.69	0.00	7.14	0.00	0.00	3.79	0.00	0.00	0.00	5.95	13.89	14.29	0.00	4.85	9.23	0.00	14.29	0.00	
llama3.1-ultrachat-fma	8B	0.00	8.51	10.17	30.00	50.00	41.43	35.72	29.55	33.33	95.24	18.18	8.33	11.11	28.57	2.94	0.00	7.69	0.00	28.57	35.71
llama3.1-tulu	8B	0.00	1.69	0.00	7.14	0.00	0.71	6.06	0.00	0.00	0.00	5.95	8.33	28.57	0.00	1.82	10.77	0.00	57.14	50.00	
llama3.1-tulu-fma	8B	0.00	4.26	11.86	30.00	35.71	42.86	36.71	27.27	36.11	98.41	18.18	10.71	11.11	42.86	2.94	2.42	6.15	0.00	35.71	50.00
qwen2.5-fma	7B	0.00	4.26	11.86	10.00	71.43	38.57	27.95	27.27	22.22	87.30	27.27	9.52	13.89	28.57	0.00	0.00	7.69	0.00	35.71	21.43
qwen2.5-coder-fma	7B	0.00	6.38	18.64	20.00	35.71	44.29	36.50	34.09	33.33	98.41	36.36	9.52	16.67	28.57	5.88	1.82	13.85	13.33	42.86	50.00
deepseek-coder-fma	7B	0.00	2.13	16.95	0.00	21.43	34.29	25.30	31.06	25.00	84.13	36.36	8.33	19.44	28.57	0.00	1.82	7.69	6.67	21.43	35.71
Pass@5																					
LLMs	Size	ProofGen						SegGen				ProofComplete				ProofInfill			Code2Sep		
		TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	ACSL	TLA	Coq	Lean	Dafny	TLA	Coq	Lean	Dafny	ACSL	ACSL
W/o Fine-tuning																					
llama3.1-instruct	8B	0.00	2.13	5.08	0.00	35.71	4.29	1.90	8.33	11.11	3.17	0.00	3.57	8.33	28.57	0.00	7.88	20.00	20.00	71.43	85.71
llama3.1-instruct	70B	7.69	8.51	13.56	20.00	35.71	27.14	6.01	17.42	13.89	52.38	36.36	11.90	27.78	28.57	0.00	10.30	30.77	20.00	35.71	57.14
qwen2.5-instruct	7B	0.00	2.13	5.08	0.00	28.57	2.86	1.39	9.09	5.56	3.17	0.00	16.67	22.22	28.57	0.00	5.45	20.00	20.00	71.43	78.57
qwen2.5-instruct	72B	15.38	6.38	22.03	30.00	21.43	17.14	3.80	14.39	25.00	15.87	9.09	16.67	38.89	28.57	0.00	11.52	26.15	33.33	57.14	92.86
qwen2.5-coder-instruct	7B	0.00	4.26	10.17	0.00	7.14	5.71	2.99	13.64	11.11	4.76	0.00	11.90	19.44	28.57	0.00	4.24	30.77	20.00	85.71	92.86
qwen2.5-coder-instruct	32B	0.00	6.38	15.25	10.00	28.57	15.71	4.96	15.15	27.78	33.33	36.36	10.71	47.22	42.86	2.94	10.91	33.85	20.00	92.86	71.43
deepseek-coder-instruct	7B	7.69	2.13	8.47	0.00	7.14	8.57	2.82	5.30	5.56	4.76	18.18	9.52	16.67	28.57	0.00	4.24	18.46	33.33	57.14	64.29
deepseek-coder-instruct	33B	0.00	6.38	6.78	20.00	7.14	2.86	3.23	6.82	13.89	25.40	0.00	11.90	19.44	42.86	0.00	10.30	40.00	26.67	28.57	14.29
star-coder-instruct	15B	15.38	4.26	10.17	30.00	50.00	34.29	3.46	7.58	27.78	0.00	45.45	16.67	41.67	57.14	0.00	16.97	38.46	26.67	42.86	42.86
W/ Fine-tuning																					
llama3.1-fma	8B	0.00	6.38	15.25	20.00	57.14	41.43	34.57	26.52	27.78	90.48	36.36	9.52	11.11	28.57	2.94	0.00	13.85	0.00	42.86	57.14
llama3.1-ultrachat	8B	0.00	3.39	0.00	14.29	0.00	0.07	7.58	0.00	1.59	0.00	4.76	22.22	14.29	0.00	4.85	9.23	6.67	14.29	0.00	
llama3.1-ultrachat-fma	8B	0.00	10.64	15.25	30.00	57.14	42.86	39.86	31.06	38.89	96.83	18.18	9.52	13.89	28.57	11.76	0.00	10.77	0.00	42.86	64.29
llama3.1-tulu	8B	0.00	2.13	1.69	0.00	21.43	0.00	1.43	6.06	0.00	3.17	0.00	13.10	30.56	28.57	0.00	6.67	21.54	13.33	78.57	64.29
llama3.1-tulu-fma	8B	0.00	6.38	11.86	30.00	35.71	44.29	40.20	32.58	38.89	98.41	18.18	13.10	11.11	42.86	2.94	4.24	9.23	0.00	71.43	64.29
qwen2.5-fma	7B	0.00	4.26	13.56	10.00	71.43	40.00	32.77	30.30	27.78	88.89	45.45	10.71	19.44	28.57	2.94	1.21	12.31	0.00	57.14	50.00
qwen2.5-coder-fma	7B	0.00	8.51	22.03	30.00	50.00	44.29	40.81	35.61	36.11	98.41	54.55	10.71	19.44	42.86	8.82	4.24	20.00	13.33	64.29	64.29
deepseek-coder-fma	7B	0.00	4.26	16.95	0.00	35.71	35.71	30.70	36.36	25.00	85.71	45.45	13.10	16.67	28.57	2.94	4.24	9.23	6.67	21.43	57.14

Table 9. Full Experiment Results on Pass@1 and Pass@5

Appendix F. Prompt Design

We listed the prompts that are used for data preparation and inference in the following. For **data preparation**, as shown in Figure 5, to generate descriptions for the given proof segments, the prompt template consists of five parts: (1) Role description, (2) Domain knowledge of TLA+, (3) Task description, (4) Few-shot examples (we show one example in the figure, while three-shots were used in RQ2), (5) The proof or proof segment to be summarized.

For the **inference**, for each task, we designed five different instructions to avoid overfitting. The prompts for each task are shown in Figure 6 – Figure 10. For each task, we first randomly choose one instruction and concat the inputs.

```
Prompt for Description Generation. Take TLA+ as an example.

# Role description
As an expert in TLA+, you are good at understanding and writing TLA+. TLA+ is a formal specification language used for modeling and verifying concurrent and distributed systems.

# Domain knowledge
1. The logical operators supported by TLA+ include:  $\wedge$  (and),  $\vee$  (or),  $\sim$  (not),  $\Rightarrow$  (Implication),  $\Leftrightarrow$  (Bidirectional implication), TRUE, FALSE,  $\forall A$  (Universal Quantification),  $\exists E$  (Existential Quantification)
2. The sets operators supported by TLA+ include: = (Equality), # (not equal),  $\cup$  (Union),  $\cap$  (Intersection),  $\in$  (Membership),  $\notin$  (Not in),  $\subseteq$  (Subset Equal),  $\setminus$  (Difference).
3. The temporal Operators supported by TLA+ include  $\square x > 0$ , which is an example of  $\square$  (Always). It means that at all times, the value of variable x is greater than 0.  $\diamond x = 0$  is an example of  $\diamond$  (Eventually). It means that at some point in time, the value of variable x becomes 0.
4. Built-in keywords and operators in TLA+ include: 'MODULE', 'EXTENDS', 'CONSTANT', 'INSTANCE', 'VARIABLE', 'ASSUME', 'PROVE', 'INIT', 'NEXT', 'ACTION', 'SPECIFICATION', 'IF', 'ELSE', 'WITH', 'CASE', 'THEN', 'LET', 'IN', 'CHOOSE', 'ENABLED', 'UNCHANGED', 'DOMAIN'.

# Task description
Given a TLA+ code snippet, you need to summarize the given TLA+ in several sentences in detail.

# Example Input and Output
## Code
Return(c,S) ==
 $\wedge S \# \{ \} \wedge S \subseteq \text{alloc}[c]$ 
 $\wedge \text{alloc}' = [\text{alloc EXCEPT } ![c] = \setminus S]$ 
 $\wedge \text{UNCHANGED } \text{unsat}$ 
## Description:
An operation 'Return(c,S)' that represents the return of a set of resources by a client. It satisfies the following conditions:
- The set 'S' is not empty and 'S' must be a subset of the set of allocated resources to the client 'c'.
- The 'alloc' is updated by assigning the difference between the current set of allocated resources and the set 'S' to the 'c' index of 'alloc' except 'c'.
- The 'unsat' remains unchanged.

# Code to be described:
<A proof segment>
```

Figure 5. Prompt for generating TLA+ description. The prompt templates for other formal specification languages are in the same structure.

Prompt for SegGen Task (lang: a placeholder to be replaced by each formal specification language name.)

Task Description (SegGen)

1. Translate the given natural language into {lang} syntax.
2. Model the intention written in natural language using {lang}.
3. Express the requirement using {lang}.
4. Model the given natural language into {lang}.
5. Translate the given requirement using {lang}'s syntax and semantics.

(Randomly choose one of the above.)

You only need to return the {lang} formal specification without explanation.

Input

<Input goes here>

Figure 6. Prompt for SegGen Task.

Prompt for ProofGen Task (lang: a placeholder to be replaced by each formal specification language name.)

Task Description (ProofGen)

1. Translate the given requirements into lang syntax.
2. Model the given requirements written in natural language using lang.
3. Express the requirements using lang.
4. Model the given requirements written in natural language into lang.
5. Translate the given requirements into lang's syntax and semantics.

(Randomly choose one of the above.)

(For ACSL): You only need to return the lang formal specification with the code without explanation.

(For others): You only need to return the lang formal specification without explanation.

Input

<Input goes here>

Figure 7. Prompt for ProofGen Task.

Prompt for ProofComplete Task (lang: a placeholder to be replaced by each formal specification language name.)

Task Description (ProofComplete)

1. Please complete the following formal proof in formal specification language lang according to the given requirement.
2. Please complete the following formal proof in lang according to the given requirement.
3. Please complete the given formal proof in lang following the requirement below.
4. Please complete the following formal proof in lang according to the requirement below.
5. Please complete the following formal proof in lang according to the given requirement.

(Randomly choose one of the above.)

You only need to return the completed lang formal specification (together with the provided formal specification) without explanation.

Input

<Input goes here>

Figure 8. Prompt for ProofComplete Task.

Prompt for ProofComplete Task (lang: a placeholder to be replaced by each formal specification language name.)

Task Description (ProofInfill)

1. Please fill in the [MASK] in the following formal proof in formal specification language lang according to the given requirement.
2. Please fill in the [MASK] in the following formal proof in lang according to the given requirement.
3. Please complete the given formal proof in lang following the requirement below by filling in the [MASK].
4. Please fill in the [MASK] in the following formal proof in lang according to the requirements below.
5. Please fill in the [MASK] in the following formal proof in lang according to the given requirement.

(Randomly choose one of the above.)

You only need to return the completed lang formal specification (together with the provided formal specification) without explanation.

Input
<Input goes here>

Figure 9. Prompt for ProofInfill Task.

Prompt for ProofComplete Task

Task Description (Code2Proof)

1. Please fill in the [MASK] in ACSL according to the given requirement and ACSL specification.
2. Please fill in the [MASK] in ACSL according to the given requirement.
3. Please fill in the [MASK] in ACSL according to the given ACSL specification.
4. Please fill in the [MASK] in ACSL according to the given requirement and ACSL specification.
5. Please infill the [MASK] in ACSL according to the given requirement.

(Randomly choose one of the above.)

You only need to return the completed ACSL formal specification (together with the provided formal specifications and C programs) without explanation.

Input
<Input goes here>

Figure 10. Prompt for Code2Proof Task.

Footnotes

¹ For ease of expression, we generally refer to *verifiable formal proofs* and *checkable models* as “proofs” for the sake of presentation simplification.

² <https://mir.fi.muni.cz/MREC/>

References

1. ^a. ^b. Yang K, Poesia G, He J, Li W, Lauter K, Chaudhuri S, Song D (2024). "Formal Mathematical Reasoning: A New Frontier in AI". arXiv preprint arXiv:2412.16075.
2. ^a. ^b. AlphaProof, DeepMind and Teams, AlphaGeometry (2024). "AI achieves silver-medal standard solving International Mathematical Olympiad problems. **25** July 2024."
3. ^a. ^b. Trinh TH, Wu Y, Le QV, He H, Luong T (2024). "Solving olympiad geometry without human demonstrations". *Nature*. 625 (7995): 476–482.
4. ^a. ^b. Li Z, Sun J, Murphy L, Su Q, Li Z, Zhang X, Yang K, Si X (2024). "A Survey on Deep Learning for Theorem Proving". arXiv preprint arXiv:2404.09939.
5. ^a. ^b. Tao T (2024). "AI Will Become Mathematicians' 'Co-Pilot'". <https://www.scientificamerican.com/article/ai-will-become-mathematicians-co-pilot/>.
6. ^a. ^b. Welleck S, Liu J, Le Bras R, Hajishirzi H, Choi Y, Cho K (2021). "NaturalProofs: Mathematical Theorem Proving in Natural Language". Preprint, arXiv:2104.0112.
7. ^a. ^b. Yang K, Swope AM, Gu A, Chalamala R, Song P, Yu S, Godil S, Prenger R, Anandkumar A (2023). "Leandojo: Theorem proving with retrieval-augmented language models". arXiv preprint arXiv:2306.15626.

8. [Welleck S, Liu J, Lu X, Hajishirzi H, Choi Y \(2022\)](#). "Naturalprover: Grounded mathematical proof generation with language models". *Advances in Neural Information Processing Systems*. 35: 4913–4927.
9. [Yang K, Deng J \(2019\)](#). "Learning to prove theorems via interacting with proof assistants." In: *International Conference on Machine Learning*. PMLR. pp. 6984–6994.
10. [Chen M, Tworek J, Jun H, Yuan Q, Ponde de Oliveira Pinto H, Kaplan J, Edwards H, Burda Y, Joseph N, Brockman G, Ray A, Puri R, Krueger G, Petrov M, Khlaaf H, Sastry G, Mishkin P, Chan B, Gray S, Ryder N, Pavlov M, Power A, Kaiser L, Bavarian M, Winter C, Tillet P, Petroski Such F, Cummings D, Plappert M, Chantzis F, Barnes E, Herbert-Voss A, Hebguss W, Nichol A, Paino A, Tezak N, Tang J, Babuschkin I, Balaji S, Jain S, Saunders W, Hesse C, Carr AN, Leike J, Achiam J, Misra V, Morikawa E, Radford A, Knight M, Brundage M, Murati M, Mayer K, Welinder P, McGrew B, Amodei D, McCandlish S, Sutskever I, Zaremba W \(2021\)](#). "Evaluating large language models trained on code". *arXiv. cs.LG: 2107.03374*.
11. [Austin J, Odena A, Nye M, Bosma M, Michalewski H, Dohan D, Jiang E, Cai C, Terry M, Le Q, et al. \(2021\)](#). "Program Synthesis with Large Language Models". *arXiv preprint arXiv:2108.07732*.
12. [Appel AW](#). "Verified Software Toolchain". In: *Barthe G, editor. Programming Languages and Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg; 2011. p. 1–17. ISBN 978–3–642–19718–5.
13. [Klein G, Elphinstone K, Heiser G, Andronick J, Cock D, Derrin P, Elkaduwe D, Engelhardt K, Kolanski R, Norrish M, Sewell T, Tuch H, Winwood S \(2009\)](#). "seL4: formal verification of an OS kernel". *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles, SOSP '09, New York, NY, USA: Association for Computing Machinery*. p. 207–220. doi:10.1145/1629575.1629596.
14. [Leroy X, Blazy S, Ke4stner D, Schommer B, Pister M, Ferdinand C \(2016\)](#). "CompCert—a formally verified optimizing compiler". In: *ERTS 2016: Embedded Real Time Software and Systems, 8th European Congress*.
15. [Hawblitzel C, Howell J, Lorch JR, Narayan A, Parno B, Zhang D, Zill B \(2014\)](#). "Ironclad apps: End-to-End security via automated Full-System verification." In: *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pages 165–181.
16. [Wang R, Zhang J, Jia Y, Pan R, Diao S, Pi R, Zhang T \(2024\)](#). "Theoremlama: Transforming general-purpose llms into lean4 experts". *arXiv preprint arXiv:2407.03203*.
17. [Wei J, Bosma M, Zhao VY, Guu K, Yu AW, Lester B, Du N, Dai AM, Le QV \(2021\)](#). "Finetuned language models are zero-shot learners". *arXiv preprint arXiv:2109.01652*.
18. [Sanh V, Webson A, Raffel C, Bach SH, Sutawika L, Alyafeai Z, Chaffin A, Stiegler A, Le Scao T, Raja A, et al. \(2021\)](#). "Multitask prompted training enables zero-shot task generalization". *arXiv preprint arXiv:2110.08207*.
19. [Tihanyi N, Jain R, Charalambous Y, Ferrag MA, Sun Y, Cordeiro LC \(2023\)](#). "A new era in software security: Towards self-healing software via large language models and formal verification". *arXiv preprint arXiv:2305.14752*.
20. [Ding N, Chen Y, Xu B, Qin Y, Hu S, Liu Z, Sun M, Zhou B \(2023\)](#). "Enhancing chat language models by scaling high-quality instructional conversations". *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Singapore: Association for Computational Linguistics. pp. 3029–3051. doi:10.18653/v1/2023.emnlp-main.183. Available from: <https://aclanthology.org/2023.emnlp-main.183/>.
21. [Wang Y, Kordi Y, Mishra S, Liu A, Smith NA, Khashabi D, Hajishirzi H \(2022\)](#). "Self-instruct: Aligning language models with self-generated instructions". *arXiv preprint arXiv:2212.10560*.
22. [Jimenez CE, Yang J, Wettig A, Yao S, Pei K, Press O, Narasimhan K \(2023\)](#). "Swe-bench: Can language models resolve real-world github issues?" *arXiv preprint arXiv:2310.06770*.
23. [Spiess C, Gros D, Pai KS, Pradel M, Rabin MRI, Alipour A, Jha S, Devanpu B, Ahmed T \(2024\)](#). "Calibration and correctness of language models for code". *arXiv preprint arXiv:2402.02047*. Available from: <https://arxiv.org/abs/2402.02047>.
24. [Tao T, Dillies Y, Mehta B \(2023\)](#). "Formalizing the proof of PFR in Lean4 using Blueprint: a short tour". *Blog post, November*.
25. [Davis AM](#). *Software requirements: analysis and specification*. Prentice Hall Press; 1990.
26. [Anton AI \(1996\)](#). "Goal-based requirements analysis." In: *Proceedings of the second international conference on requirements engineering*. IEEE. pp. 136–144.
27. [Grady JO](#). *System requirements analysis*. Elsevier; 2010.
28. [Jin Z](#). *Environment modeling-based requirements engineering for software intensive systems*. Morgan Kaufmann; 2017.
29. [Taggart Jr W, Tharp MO \(1977\)](#). "A survey of information requirements analysis techniques". *ACM Computing Surveys (CSUR)*. 9 (4): 273–290.
30. [Deeptimahanti DK, Babar MA \(2009\)](#). "An automated tool for generating UML models from natural language requirements." In: *2009 IEEE/ACM International Conference on Automated Software Engineering*. IEEE. pp. 680–682.

31. [△]Javed M, Lin Y (2021). "iMER: Iterative process of entity relationship and business process model extraction from the requirements". *Information and Software Technology*. 135: 106558.
32. [△]Wang Y, Chen JW, Xia X, Jiang B (2021). "Intelligent requirements elicitation and modeling: A literature review". *Journal of Computer Research and Development*. 58 (4): 683–705.
33. [△]Bjin D, Zhao S, Jin Z, Chen X, Wang C, Fang Z, Xiao H (2024). "An Evaluation of Requirements Modeling for Cyber-Physical Systems via LLMs". arXiv preprint arXiv:2408.02450.
34. [△]Zhou Q, Li T, Wang Y (2022). "Assisting in requirements goal modeling: a hybrid approach based on machine learning and logical reasoning." In: *Proceedings of the 25th International Conference on Model Driven Engineering Languages and Systems*, pp. 199–209.
35. [△]Fatwanto A (2012). "Translating software requirements from natural language to formal specification." In: *2012 IEEE International Conference on Computational Intelligence and Cybernetics (CyberneticsCom)*. IEEE. pp. 148–152.
36. [△]Davril JM, Delfosse E, Hariri N, Acher M, Cleland-Huang J, Heymans P (2013). "Feature model extraction from large collections of informal product descriptions". *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2013, New York, NY, USA: Association for Computing Machinery*. p. 290–300. doi:10.1145/2491411.2491455.
37. [△]Wu Y, Jiang AQ, Li W, Rabe M, Staats C, Jamnik M, Szegedy C (2022). "Autoformalization with large language models". *Advances in Neural Information Processing Systems*. 35: 32353–32368.
38. [△]Jiang AQ, Welleck S, Zhou JP, Li W, Liu J, Jamnik M, Lacroix T, Wu Y, Lample G (2022). "Draft, sketch, and prove: Guiding formal theorem provers with informal proofs". arXiv preprint arXiv:2210.12283.
39. [△]Chaudhary S (2023). "Code Alpaca: An Instruction-following LLaMA model for code generation". GitHub repository. Available from: <https://github.com/sahil28014/codealpaca>.
40. [△]Sun C, Sheng Y, Padon O, Barrett C (2024). "Clover: Closed-loop verifiable code generation". Preprint, arXiv:2310.17807. Available from: <https://arxiv.org/abs/2310.17807>.
41. [△]Raychev V, Vechev M, Yahav E (2014). "Code completion with statistical language models". In: *Proceedings of the 35th ACM SIGPLAN conference on programming language design and implementation*. pp. 419–428.
42. [△]Husein RA, Aburajouh H, Catal C (2024). "Large language models for code completion: A systematic literature review". *Computer Standards & Interfaces*. Page 103917.
43. [△]Svyatkovskiy A, Zhao Y, Fu S, Sundaresan N (2019). "Pythia: Ai-assisted code completion system." In: *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2727–2735.
44. [△]Dakhel AM, Majdinasab V, Nikanjam A, Khomh F, Desmarais MC, Jiang ZMJ (2023). "Github copilot ai pair programmer: Asset or liability?" *Journal of Systems and Software*. 203: 111734.
45. [△]Song P, Yang K, Anandkumar A (2024). "Towards large language models as copilots for theorem proving in lean". arXiv preprint arXiv:2404.12534.
46. [△]Fried D, Aghajanyan A, Lin J, Wang S, Wallace E, Shi F, Zhong R, Yih W, Zettlemoyer L, Lewis M (2022). "InCoder: A generative model for code infilling and synthesis". arXiv preprint arXiv:2204.05999.
47. [△]Baudin P, Filliatre JC, Marche C, Monate B, Moy Y, Prevosto V (2021). "ACSL: ANSI/ISO C Specification". Citeseer.
48. [△]Hatcliff J, Leavens GT, Leino KRM, M\uoofcller P, Parkinson M (2012). "Behavioral interface specification languages". *ACM Comput. Surv.* 44 (3): Article 16, 58 pages.
49. [△]Wen C, Cao J, Su J, Xu Z, Qin S, He M, Li H, Cheung S, Tian C (2024). "Enchanting program specification synthesis by large language models using static analysis and program verification." In: *International Conference on Computer Aided Verification*. Springer. pp. 302–328.
50. [△]Ma L, Liu S, Li Y, Xie X, Bu L (2024). "SpecGen: Automated Generation of Formal Program Specifications via Large Language Models". Preprint, arXiv:2401.08807.
51. [△]Huet G. *The calculus of constructions*. PhD thesis. INRIA; 1986.
52. [△]Leino KRM (2010). "Dafny: An automatic program verifier for functional correctness." In: *International conference on logic for programming artificial intelligence and reasoning*. Springer. pp. 348–370.
53. [△]Moura L de, Ullrich S (2021). "The Lean 4 Theorem Prover and Programming Language". In *Automated Deduction – CADE 28: 28th International Conference on Automated Deduction, Virtual Event, July 12–15, 2021, Proceedings, Berlin, Heidelberg: Springer-Verlag*. p. 625–635. doi:10.1007/978-3-030-79876-5_37.

54. ^{ab}Yu Y, Manolios P, Lamport L. Model checking TLA+ specifications. In: *Advanced research working conference on correct hardware design and verification methods*. Springer; 1999. p. 54–66.
55. ^ΔLamport L (2002). "Specifying systems: the TLA+ language and tools for hardware and software engineers."
56. ^ΔAvigad J, Buzzard K, Lewis RY, Massot P (2020). "Mathematics in Lean". Technical report, Lean community. En <https://leanprover-community>.
57. ^{ab}Cuoq P, Kirchner F, Kosmatov N, Prevosto V, Signoles J, Yakobowski B (2012). "Frama-C: A software analysis perspective." In: *International conference on software engineering and formal methods*. Springer. pp. 233–247.
58. ^ΔJhala R, Majumdar R (2009). "Software model checking". *ACM Computing Surveys (CSUR)*. 41 (4): 1–54.
59. ^ΔClarke EM. Model checking. In: *Foundations of Software Technology and Theoretical Computer Science: 17th Conference Kharagpur, India, December 18–20, 1997 Proceedings 17*. Springer; 1997. p. 54–56.
60. ^ΔGPT-4 and GPT-4-turbo Preview (2023). Available from: <https://learn.microsoft.com/en-us/azure/ai-services/openai/concepts/models#-gpt-4-and-gpt-4-turbo-preview>.
61. ^{ab}Wang Y, Ivison H, Dasigi P, Hessel J, Khot T, Chandu K, Wadden D, MacMillan K, Smith NA, Beltagy I, et al. (2023). "How far can camels go? exploring the state of instruction tuning on open resources". *Advances in Neural Information Processing Systems*. 36: 74764–74786.
62. ^ΔLu Y, Liu J, Zhang Y, Liu Y, Tian X (2022). "Prompt distribution learning". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5206–5215.
63. ^ΔFeng CM, Yu K, Liu Y, Khan S, Zuo W (2023). "Diverse data augmentation with diffusions for effective test-time prompt tuning." In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2704–2714.
64. ^ΔSanh V, Webson A, Raffel C, Bach SH, Sutawika L, Alyafeai Z, Chaffin A, Stiegler A, Le Scao T, Raja A, Dey M, Bari MS, Xu C, Thakker U, Sharma SS, Szczechla E, Kim T, Chhablani G, Nayak N, Datta D, Chang J, Jiang MTJ, Wang H, Manica M, Shen S, Yong ZX, Pandey H, Bawden R, Wang T, Neeraj T, Rozen J, Sharma A, Santilli A, Fevry T, Fries JA, Teehan R, Bers T, Biderman S, Gao L, Wolf T, Rush AM (2022). "Multitask prompted training enables zero-shot task generalization". Preprint, arXiv:2110.08207.
65. ^ΔYu Y, Kuppe M. TLC Model Checker. Available from: <https://tla.msr-inria.inria.fr/tlatoolbox/doc/model/executing-tlc.html>.
66. ^ΔCarvalho N, da Silva Sousa C, Pinto JS, Tomb A. Formal Verification of kLBCB with the WP Frama-C Plug-in. In: *NASA Formal Methods: 6th International Symposium, NFM 2014, Houston, TX, USA, April 29–May 1, 2014. Proceedings 6*. Springer; 2014. p. 343–358.
67. ^ΔPapineni K, Roukos S, Ward T, Zhu WJ. (2002). "Bleu: a Method for Automatic Evaluation of Machine Translation". In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, Philadelphia, Pennsylvania, USA: Association for Computational Linguistics*; p. 311–318. doi:10.3115/1073083.1073135. Available from: <https://aclanthology.org/P02-1040/>.
68. ^{ab}Meta AI (2024). "Introducing meta llama 3: The most capable openly available llm to date". Blog. <https://ai.meta.com/blog/meta-llama-3/>. Online; accessed 15-January-2024.
69. ^{ab}Yang A, Yang B, Zhang B, Hui B, Zheng B, Yu B, Li C, Liu D, Huang F, Wei H, et al. (2024). "Qwen2.5 Technical Report". arXiv preprint arXiv:2412.15115.
70. ^ΔHui B, Yang J, Cui Z, Yang J, Liu D, Zhang L, Liu T, Zhang J, Yu B, Lu K, et al. (2024). "Qwen2. 5-coder technical report". arXiv preprint arXiv:2409.12186.
71. ^ΔLozhkov A, Li R, Ben Allal L, Cassano F, Lamy-Poirier J, Tazi N, Tang A, Pykhtar D, Liu J, Wei Y, et al. (2024). "StarCoder 2 and the stack v2: The next generation". arXiv preprint arXiv:2402.19173. Available from: <https://arxiv.org/abs/2402.19173>.
72. ^{ab}Guo D, Zhu Q, Yang D, Xie Z, Dong K, Zhang W, Chen G, Bi X, Wu Y, Li YK, et al. (2024). "DeepSeek-Coder: When the Large Language Model Meets Programming--The Rise of Code Intelligence". arXiv preprint arXiv:2401.14196.
73. ^ΔGuo D, Yang D, Zhang H, Song J, Zhang R, Xu R, Zhu Q, Ma S, Wang P, Bi X, et al. (2025). "Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning". arXiv preprint arXiv:2501.12948.
74. ^{ab}Ivison H, Wang Y, Pyatkin V, Lambert N, Peters M, Dasigi P, Jang J, Wadden D, Smith NA, Beltagy I, et al. (2023). "Camels in a changing climate: Enhancing llm adaptation with tulu 2". arXiv preprint arXiv:2311.10702.
75. ^{ab}Lambert N, Morrison J, Pyatkin V, Huang S, Ivison H, Brahman F, Miranda LJV, Liu A, Dziri N, Lyu S, et al. 2024. "Tulu 3: Pushing frontiers in open language model post-training". arXiv preprint arXiv:2411.15124.
76. ^ΔHendrycks D, Burns C, Kadavath S, Arora A, Basart S, Tang E, Song D, Steinhardt J (2021). "Measuring Mathematical Problem Solving With the MATH Dataset". *NeurIPS*.
77. ^ΔCobbe K, Kosaraju V, Bavarian M, Chen M, Jun H, Kaiser L, Plappert M, Tworek J, Hilton J, Nakano R, Hesse C, Schulman J (2021). "Training Verifiers to Solve Math Word Problems". arXiv preprint arXiv:2110.14168.

78. ^ΔSuzgun M, Scales N, Schärli N, Gehrmann S, Tay Y, Chung HW, Chowdhery A, Le QV, Chi EH, Zhou D, Wei J (2022). "Challenging BIG-Bench Tasks and Whether Chain-of-Thought Can Solve Them". arXiv preprint arXiv:2210.09261.
79. ^ΔAustin J, Odena A, Nye M, Bosma M, Michalewski H, Dohan D, Jiang E, Cai C, Terry M, Le Q, et al. (2021). "Program Synthesis with Large Language Models". arXiv preprint arXiv:2108.07732.
80. ^ΔHoltzman A, Buys J, Du L, Forbes M, Choi Y (2020). "The Curious Case of Neural Text Degeneration". In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net.
81. ^ΔKwon W, Li Z, Zhuang S, Sheng Y, Zheng L, Yu CH, Gonzalez JE, Zhang H, Stoica I (2023). "Efficient Memory Management for Large Language Model Serving with PagedAttention". Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles.
82. ^ΔCook SA (2023). "The complexity of theorem-proving procedures." In: Logic, automata, and computational complexity: The works of Stephen A. Cook, pp. 143-152.
83. ^ΔFerreira D, Freitas A (2020). "Natural language premise selection: Finding supporting statements for mathematical text". arXiv preprint arXiv:2004.14959.
84. ^ΔLi WW, Ziser Y, Coavoux M, Cohen SB (2023). "BERT is not the count: Learning to match mathematical statements with proofs". arXiv preprint arXiv:2302.09350.
85. ^ΔThe mathlib Community (2020). "The lean mathematical library". In Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, POPL'20. ACM. doi:10.1145/3372885.3373824.
86. ^ΔZheng K, Han JM, Polu S (2021). "MiniF2F: a cross-system benchmark for formal Olympiad-level mathematics". arXiv preprint arXiv:2109.00110.
87. ^ΔLiu C, Shen J, Xin H, Liu Z, Yuan Y, Wang H, Ju W, Zheng C, Yin Y, Li L, Zhang M, Liu Q (2023). "FIMO: A Challenge Formal Dataset for Automated Theorem Proving". Preprint, arXiv: 2309.04295.
88. ^ΔAzerbayev Z, Piotrowski B, Schoelkopf H, Ayers EW, Radev D, Avigad J (2023). "ProofNet: Autoformalizing and Formally Proving Undergraduate-Level Mathematics". Preprint, arXiv:2302.12433. Available from: <https://arxiv.org/abs/2302.12433>.
89. ^ΔMugnier E, Anaya Gonzalez E, Jhala R, Polikarpova N, Zhou Y (2024). "Laurel: Generating Dafny Assertions Using Large Language Models". Preprint, arXiv:2405.16792. Available from: <https://arxiv.org/abs/2405.16792>.
90. ^ΔDella Penna G, Intrigila B, Magazzeni D, Melatti I, Tronci E (2013). "CgMurphi: Automatic synthesis of numerical controllers for nonlinear hybrid systems". European Journal of Control. 19 (1): 14-36.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.