PROVER AGENT: AN AGENT-BASED FRAMEWORK FOR FORMAL MATHEMATICAL PROOFS

Anonymous authors

Paper under double-blind review

ABSTRACT

We present Prover Agent, a novel AI agent for automated theorem proving that integrates large language models (LLMs) with a formal proof assistant, Lean. Prover Agent coordinates an informal reasoning LLM, a formal prover model, and feedback from Lean while also generating auxiliary lemmas to assist in discovering the overall proof strategy. It achieves an 88.1% success rate on the MiniF2F benchmark, establishing a new state-of-the-art among methods using small language models (SLMs) with a much lower sample budget than previous approaches. We also present theoretical analyses and case studies that illustrate how these generated lemmas contribute to solving challenging problems.

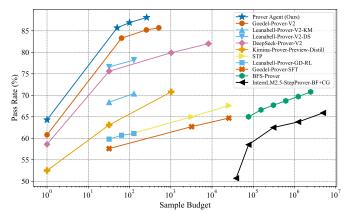


Figure 1: Comparison of theorem-proving performance on the MiniF2F benchmark (Zheng et al., 2022) among methods using SLMs. Our approach achieves a higher success rate with fewer sample budgets, establishing a new state-of-the-art at this scale.

1 Introduction

Recent advances in the reasoning capabilities of large language models (LLMs) have driven remarkable progress across many areas of artificial intelligence, including mathematical theorem proving and problem solving (OpenAI, 2024; DeepSeek-AI, 2025; Yang et al., 2025a; Lewkowycz et al., 2022). However, LLMs are prone to errors and hallucinations that can undermine their reliability (Ji et al., 2023; Huang et al., 2025; Xu et al., 2025). Inference-time scaling techniques such as chain-of-thought have greatly enhanced their reasoning performance by allowing models to reflect on and correct faulty reasoning steps (OpenAI, 2024; DeepSeek-AI, 2025; Wei et al., 2022). Nonetheless, eliminating mistakes entirely remains challenging, especially for more difficult problems (Wei et al., 2022; Zeng et al., 2025).

Formal proof assistants such as Lean (Moura & Ullrich, 2021), The Rocq Prover (previously known as Coq) (Barras et al., 1999), and Isabelle (Paulson, 1994) rigorously verify by computer that every inference step in mathematical proofs written in their respective languages is correct, based on the Curry–Howard correspondence. This helps mathematicians verify the correctness of proofs. Here, no errors, omissions of detail, implicit assumptions, or ambiguities are permitted. However, working

with formal proof assistants typically requires painstaking manual effort and meticulous detail. As a result, automating mathematical theorem proving has long been a grand challenge in artificial intelligence and formal methods (Newell & Simon, 1956; Irving et al., 2016; Polu & Sutskever, 2020; Jiang et al., 2023; Lu et al., 2023).

Consequently, formal theorem proving with LLMs has become increasingly important in recent years, leading to a growing body of research in this area (Wang et al., 2024b; Wu et al., 2024a; Xin et al., 2025b; Li et al., 2025; Xin et al., 2025a; Dong & Ma, 2025; Lin et al., 2025b; Zhang et al., 2025; Wang et al., 2025; Ren et al., 2025). This not only provides a way to guarantee the correctness of mathematical reasoning by LLMs, but also marks a major breakthrough in automated theorem proving. A key point is the complementary strengths of LLMs and formal proof assistants: LLMs excel in reasoning and generation but may produce errors and lack guarantees of correctness, whereas formal proof assistants, such as Lean, possess perfect verification capabilities grounded in mathematical logic but are not generative.

Yet, significant hurdles remain in bridging informal reasoning and formal proving (Yang et al., 2025b). For instance, prompting o3-mini (OpenAI, 2025) to directly generate a complete Lean proof for a competition-level problem succeeds in only 6.0% of cases in a single attempt, despite its strong performance on competition-level mathematical reasoning in natural language (Yousefzadeh & Cao, 2025). Even when fine-tuned on mathematical data, trained with reinforcement learning, or allowed chain-of-thought, purely neural approaches fail to produce correct formal proofs, and their formal proving capabilities still lag far behind their informal reasoning skills in natural language.

To bridge this gap between informal reasoning and formal proving, we propose a novel agent framework (**Prover Agent**) that coordinates an informal reasoning LLM, a formal prover model, and the Lean verification system. To tackle difficult problems that cannot be solved directly, the agent generates auxiliary lemmas to assist in discovering a viable proof strategy. These lemmas are not limited to subgoals that can be directly inserted into a formal proof, but may also include special cases or intermediate facts derived from the assumptions. Such lemmas are particularly useful when the overall proof strategy is not apparent from the outset, as they help in constructing a viable plan. On the MiniF2F benchmark (Zheng et al., 2022), it achieves an 88.1% success rate, establishing a new state-of-the-art among methods using small language models (SLMs). Notably, it uses only SLMs with much smaller sample budget than previous high-performing approaches, making it much more efficient in terms of inference-time cost. Furthermore, we provide both a theoretical analysis and a case study to demonstrate the effectiveness of our agent's approach to generating auxiliary lemmas.

Our contributions are summarized as follows:

- Coordination of Informal and Formal Reasoning with Lean Feedback: Our agent combines an informal LLM and a formal prover under Lean's verification. The LLM produces natural language reasoning and lemmas, which the prover formalizes and Lean checks. Errors detected by Lean are immediately fed back, enabling iterative refinement of constructed proofs.
- Auxiliary Lemma Generation for Strategy Discovery: For challenging problems that cannot be solved directly, our agent generates auxiliary lemmas, such as specific cases, intermediate facts, or hypothesis-driven conjectures, which are then formally proved. By reconsidering the overall proof in light of the verified lemmas, the system uncovers viable proof strategies even when the solution path is not apparent at first.
- State-of-the-Art Theorem-Proving Performance: On the challenging MiniF2F benchmark (Zheng et al., 2022), a standard benchmark for formal theorem proving that consists of 488 problems drawn from mathematics Olympiads and advanced mathematics, our agent achieves 88.1% pass rate, establishing a new state-of-the-art among methods using SLMs.
- Efficiency in Inference-Time Cost: The 88.1% success rate was achieved using only SLMs with a much smaller sample budget than previous high-performing approaches. This emphasizes the efficiency of our approach in terms of inference-time cost.

2 RELATED WORK

In this section, we provide a brief overview of recent advancements in automated formal theorem proving. Details of representative systems are provided in Appendix A.

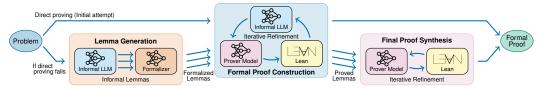


Figure 2: Overall workflow of Prover Agent. The agent coordinates informal reasoning, formal proving, and Lean verification. It first attempts direct proving; if unsuccessful, it generates auxiliary lemmas to guide the discovery of a viable proof strategy. These lemmas are then formally proved, and the successfully proved lemmas are subsequently used to synthesize the final proof.

Tree-Search-based Formal Proving. Tree-search methods construct Lean proofs tactic-by-tactic and navigate the proof space with explicit search, such as best-first search or Monte-Carlo tree search (MCTS) (Lample et al., 2022; Wang et al., 2023; Wu et al., 2024a; Zhou et al., 2024; Li et al., 2025; Xin et al., 2025a;b). This line began with stepwise tactic prediction guided by a goal state, and matured into systems that jointly optimize (i) the tactic policy, (ii) the search heuristic, and (iii) data curation for longer proofs.

Whole-Proof Generation. A complementary line to tree-search methods is whole-proof generation (First et al., 2023), where a model emits an entire Lean script in one shot, often accompanied by a long chain-of-thought reasoning trace. This approach has progressed via expert-iteration pipelines that recycle verified proofs back into training (Polu et al., 2023; Wu et al., 2021; 2024a; Lin et al., 2025a; Dong & Ma, 2025; Lin et al., 2025b;c) and via reinforcement learning with formal verifier feedback (Kaliszyk et al., 2018; Xin et al., 2025a; Zhang et al., 2025; Wang et al., 2025; Ren et al., 2025; Gloeckle et al., 2024; Ji et al., 2025; Lin et al., 2025c).

Formal Theorem Proving with Retrieval-Augmented Generation. Another emerging direction is to combine LLM-based provers with retrieval-augmented generation (RAG), where external knowledge sources or proof libraries are queried at inference time to supplement the model's reasoning (Yang et al., 2023; Shen et al., 2025)

Proof Refinement and Subgoal Decomposition. Some work has explored proof refinement, where an initial proof attempt is improved based on feedback from the proof assistant (Thakur et al., 2024; Zhou et al., 2025; Chen et al., 2025; Lin et al., 2025c). Another line of work involves subgoal decomposition, where a complex theorem is broken down into simpler subgoals that are easier to prove (Dong et al., 2025; Wang et al., 2024a; Ren et al., 2025; Zhou et al., 2025), often guided by natural-language sketches (Jiang et al., 2023; Cao et al., 2025).

The subgoal decomposition approach shares certain similarities with ours, but our method adopts a more comprehensive strategy that subsumes it. In these works, the full sketch of the proof must be correctly envisioned upfront, which is often challenging. In contrast, our approach does not assume that the overall proof strategy is fully visible from the beginning. Rather than limiting decomposition to subgoals directly aligned with a pre-defined proof plan, we also consider auxiliary lemmas, such as specific cases or intermediate facts to help developing a strategy in a bottom-up manner.

3 Method

The overall workflow is illustrated in Figure 2 and the corresponding pseudocode is shown in Algorithm 1. Given a formal math problem, our agent first attempts a direct proof, which is often sufficient for simpler problems. For more difficult problems that cannot be solved directly, it generates auxiliary lemmas to uncover a viable proof strategy. These lemmas are then formalized and proved individually, and the resulting proven lemmas are used to synthesize a final proof of the original problem. Throughout this process, feedback from Lean is used to iteratively refine constructed proofs. We describe each stage below, highlighting how the informal LLM, formal prover model, and Lean coordinate to construct formal proofs.

3.1 FORMAL PROOF CONSTRUCTION GUIDED BY INFORMAL REASONING AND ITERATIVE FEEDBACK

The agent first attempts to directly prove the given problem or a generated lemma without decomposition. To leverage the stronger mathematical reasoning ability of the informal LLM compared to that of the formal prover model, we first generate an informal proof in natural language for the given problem or lemma using the informal LLM. The formal prover model then uses the informal proof as contextual guidance to generate a formal proof, which is subsequently verified by Lean. If the proof is successful, this step is complete. If the proof fails, these steps are repeated until a successful proof is found or the maximum number of attempts $N_{\rm init}$ is reached. This process helps establish a better initial outline for the subsequent iterative refinement process.

If the proof still fails, the agent enters an iterative refinement stage. The proof with the fewest Lean verification errors among the prior attempts is selected as the initial draft. This proof is then iteratively refined based on the feedback from Lean. In each iteration, the previous proof attempt, along with the error locations and corresponding error messages, is provided to the prover model, which revises and generates a corrected version of the proof. This process is repeated until the proof is successfully verified by Lean or the maximum number of attempts $N_{\rm refine}$ is reached.

This iterative refinement process leverages Lean's verification to identify and correct mistakes. It serves as a form of self-correction through in-context learning, akin to how humans improve their understanding from feedback. This provides an efficient remedy to a key limitation of inference-time scaling with chain-of-thought, where simply increasing the number of reasoning steps does not guarantee better results due to the model's limited ability of self-correction (Zeng et al., 2025; Song et al., 2025; Stechly et al., 2025; Huang et al., 2024).

It is accessible if a generated lemma cannot be proven. This mirrors how human mathematicians often approach problems: when the overall strategy is unclear at the beginning, they may explore several directions, some of which turn out to be unproductive and are eventually discarded in favor of more promising ones. Alternatively, to handle cases where the lemma is still too challenging to prove, the system may recursively introduce smaller auxiliary lemmas, up to a depth limit D.

3.2 Lemma Generation via Informal Reasoning

When the direct proving approach fails to solve the problem, the agent generates several auxiliary lemmas. These are not limited to subgoals that can be directly inserted into a final proof; they may also include specific cases or potentially useful facts derived from the assumptions that help in developing a proof strategy. This represents a key difference from prior work, which typically relies on decomposing the problem into subgoals based on a pre-defined proof sketch (Jiang et al., 2023; Wang et al., 2024a; Ren et al., 2025; Cao et al., 2025; Zhou et al., 2025). In such approaches, it is necessary to come up with the correct overall proof strategy beforehand, which is often a challenging task. Indeed, these methods often rely on larger, stronger models such as DeepSeek-V3 (DeepSeek-AI, 2024) and DeepSeek-R1 (DeepSeek-AI, 2025) to accurately predict the entire proof plan from the outset. In contrast, our approach does not assume that the proof strategy is visible from the outset. Instead, by generating auxiliary lemmas, the agent can gradually construct an effective proof strategy in a bottom-up manner, even when the full structure is not initially apparent.

For example, when trying to prove that $n^2 + an$ is even for a natural number n and an odd number a, it may be helpful to first consider specific cases such as a=1 or a=3, i.e., n^2+n or n^2+3n . These specific cases can help reveal patterns and guide the overall proof strategy for n^2+an , even though expressions like n^2+n or n^2+3n may not explicitly appear as steps within the final proof.

This approach mirrors how human mathematicians typically work. When the overall strategy is not clear at the beginning, they often explore specific cases or consider what can be derived from the assumptions. Through such trial and error, they gradually discover the overall proof strategy.

The system first generates lemmas in natural language to leverage the stronger mathematical reasoning capabilities of the informal LLM. These lemmas are then converted into formal statements by a formalizer model, which formalizes only their assumptions and conclusions with no proof attempt. Lean is also used here to verify the syntactic correctness of the formalized statements, which are regenerated until they become valid. These formally stated lemmas are then proved using the proof construction process described in Section 3.1.

3.3 Final Proof Synthesis Guided by Verified Lemmas and Iterative Feedback

After attempting to prove each of these lemmas individually, the agent reconsiders the overall proof. With the verified lemmas as context, it attempts to to construct a proof up to $N_{\rm init}$ times, followed by iterative refinement for up to $N_{\rm refine}$ attempts, as described in Section 3.1.

4 THEORETICAL ANALYSIS

We present theoretical analyses to justify the effectiveness of our approach described in Section 3. The use of lemmas serves two key purposes: (i) decomposing proof steps under a given strategy to make them more manageable, and (ii) helping discover proof strategies when the appropriate one is not initially clear (e.g., by testing specific cases). Prior work has largely focused only on (i), often requiring larger models to directly devise an overall strategy (Wang et al., 2024a; Jiang et al., 2023; Ren et al., 2025; Cao et al., 2025; Zhou et al., 2025), whereas our approach leverages both (i) and (ii) to solve difficult problems more effectively. Sections 4.1 and 4.2 present brief results of theoretical analyses on lemma usage in cases (i) and (ii), respectively. See Appendix C for the details.

4.1 Benefits of Lemmas for Structured Proof Decomposition

Assumption 4.1. For a certain class of theorems, it is necessary to satisfy m essential intermediate facts F_1, \ldots, F_m .

Assumption 4.2. The probability p_i that the model correctly produces each F_i in a single attempt is independent across i within one global generation.

Assumption 4.3. Given a set of completed intermediate facts $\{F_i\}_{i\in S}$ with $S\subseteq [m]^1$, the probability of proving their composition F_S (e.g., simply concatenating them) is higher than the probability of proving F_S without being given those facts: $\mathbb{P}(F_S \mid \{F_i\}_{i\in S}) > \mathbb{P}(F_S)$.

Assuming $p=p_1=\cdots=p_m$ for simplicity, the following theorems hold. Rigorous versions without this simplification and without asymptotic notation are provided in Appendix C.1.

Theorem 4.4 (Required Number of Trials). Let N_{dir} denote the number of trials required to directly prove a problem T with probability at least $1-\delta$. Let N_{lem} denote the total number of trials required to complete the proof of T with probability at least $1-\delta$, when lemmas L_1,\ldots,L_n are introduced with an allowed failure probability δ_{lem} . Suppose each lemma L_i contains a subset of the essential intermediate facts $\{F_i\}_{i\in S_i}$ with $S_i\subseteq [m]$. Then the following holds:

$$N_{\text{dir}} = \Theta(p^{-m}), \qquad \mathbb{E}[N_{\text{lem}}] = \tilde{\Theta}(p^{-s}),$$

where $s := \max\{\max_i |S_i|, |R_0|\} \le m$, $R_0 := [m] \setminus \bigcup_{i=1}^n S_i$, and $\tilde{\Theta}$ indicates asymptotic order ignoring higher-order terms in δ_{lem} , which vanish when δ_{lem} is sufficiently small.

Theorem 4.5 (Threshold Condition for Lemma Efficiency). There exists a threshold $\tau \in [0,1]$ such that if $p \leq \tau$, then $\mathbb{E}[N_{\mathrm{lem}}] \leq N_{\mathrm{dir}}$ holds for any $\delta, \delta_{\mathrm{lem}} \in (0,1)$.

Theorem 4.6 (Optimal Partition of Lemma Coverage). Under the fixed lemma coverage $U := \bigcup_{i=1}^n S_i \subseteq [m]$, $\mathbb{E}[N_{\text{lem}}]$ is minimized when $|S_i| = \lceil |U|/n \rceil$ or $\lfloor |U|/n \rfloor$ for all $i \in [n]$.

The proofs are provided in Appendix C.1. Theorem 4.4 shows that lemma-based decomposition yields an exponential improvement in the order of required trials, while Theorem 4.5 indicates that for small p (i.e., difficult problems), lemma usage reduces the required number of trials. This justifies our approach of generating lemmas for difficult problems while solving easy ones directly. Furthermore, Theorem 4.6 suggests that the optimal lemmas are those that divide the problem into subproblems of approximately equal difficulty.

4.2 BENEFITS OF LEMMAS FOR DISCOVERING PROOF STRATEGIES (E.G., SPECIFIC CASES)

Let S be the set of possible proof strategies (e.g., induction, bounding with monotonicity, or case analysis with known results). Let π_0 denote the prior distribution over strategies that the model possesses, from which a strategy is chosen in the absence of any additional information. Our agent

[[]m] denotes the set $\{1, 2, \ldots, m\}$.

conducts experiments with lemmas L_1, \ldots, L_n and verifies them in Lean, thereby obtaining observations Y_1, \ldots, Y_n . By incorporating these observations into the context, the distribution is updated to the posterior $\pi_n(\cdot) := \pi(\cdot \mid Y_{1:n})$, where $Y_{1:n} := \{Y_1, \ldots, Y_n\}$, aiming to increase the probability of selecting the correct proof strategy.

Let $p(z) \in [0,1]$ denote the model's success probability under a given strategy $z \in \mathcal{S}$, and define $r \coloneqq \inf_z p(z)$. As shown in Section 4.1, this quantity can be increased by using decomposition-type lemmas. Define the entropy of the prior distribution as $H_0 \coloneqq H(Z) = -\sum_{z \in \mathcal{S}} \pi_0(z) \log \pi_0(z)$.

Theorem 4.7 (Success Probability Improvement by Lemmas). The success probability of performing one trial of final proving by sampling a strategy from the posterior distribution π_n is bounded as: $\mathbb{E}[\mathbb{P}(\text{succ}@1)] \geq r \exp(-H_0 + I(Z; Y_{1:n})).$

The proof is provided in Appendix C.2. This shows that the success probability improves exponentially in the mutual information contributed by lemmas, $I(Z; Y_{1:n})$. In particular, it exceeds the no-lemma case, where $I(Z; Y_{1:n}) = 0$.

Furthermore, this result implies that not only lemmas but any information in the context that shares mutual information with the final correct proof can similarly improve the success probability, thereby justifying our use of natural language proofs and Lean feedback.

5 EXPERIMENTS

5.1 EXPERIMENTAL SETUP

We evaluate our approach on the MiniF2F benchmark (Zheng et al., 2022), a standard dataset for formal theorem-proving systems. We use <code>DeepSeek-R1-0528-Qwen3-8B</code> (DeepSeek-AI, 2025) for the informal reasoning LLM and <code>DeepSeek-Prover-V2-7B</code> (Ren et al., 2025) and <code>Goedel-Prover-V2-8B</code> (Lin et al., 2025c) for the prover model. We set $N_{\rm init}=N_{\rm refine}=50$. Thus, the sample budget at the initial direct proving stage is 50 at the first iteration, and 100 in total when including iterative refinement. For lemmas, we use $N_{\rm init}=N_{\rm refine}=10$ for each of the three lemmas. In the final synthesis stage, $N_{\rm init}=N_{\rm refine}=50$ is used again, resulting in a total sample budget of $50+50+(10+10)\times3+50+50=260$.

There are several bugs that may result in invalid Lean proofs being incorrectly accepted, such as the user-interference bug related to the apply? tactic discussed in Ren et al. (2025), and a bug in REPL². To avoid these issues and prevent invalid proofs from being mistakenly judged as correct, we check proofs with lake build instead of REPL and additionally verified that the apply? tactic is not used. Also, to avoid this bug and obtain reliable baseline results, we re-ran the experiments for Goedel-Prover-V2-8B. We used the official prompts provided on GitHub³ and Hugging Face⁶, while keeping all other experimental settings strictly identical to those used in our method, thereby ensuring a fair comparison. For DeepSeek-Prover-V2, we relied on the results reported in (Ren et al., 2025), in which this bug has been fixed. See Appendix D for further details.

5.2 Main Result: Comparison with the Previous State-of-the-Art

The results are shown in Table 1 and Figure 3. Our agent achieves an 88.1% success rate, establishing a new state-of-the-art among methods using small language models (SLMs). Note that our agent achieves this result with a sample budget of only 260, far smaller than that of prior work, highlighting its efficiency in inference-time cost.

5.3 MODULAR AND SCALABLE DESIGN

To demonstrate the robustness of our approach, we conduct experiments across several models, namely DeepSeek-Prover-V2 and Goedel-Prover-V2. In both settings, our approach achieves higher success rates with a smaller sample budget than the vanilla versions of these models, as shown in Table 1. Furthermore, our approach can also ensemble these models. In experiments where the sample budget is split evenly between them, our agent achieves an even higher success rate, where the models complement each other on problems that one alone cannot solve. Unlike monolithic

²https://github.com/leanprover-community/repl/issues/44

³https://github.com/Goedel-LM/Goedel-Prover-V2

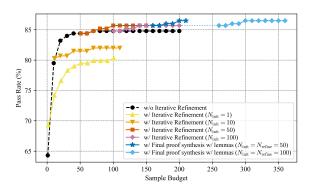
Table 1: Comparison of formal theorem-proving performance on miniF2F-test. The results are reported as the percentage of theorems proved correctly. For Prover Agent, sample budget includes all proof attempts across the full pipeline, including initial direct proving, iterative refinement, lemma proving, and final proof synthesis. The best results within each model scale are highlighted in **bold**.

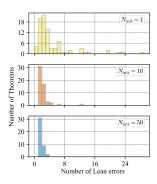
Prover System			Method	Model Size	Sample Budget	miniF2F-test
Large Language Models						
DSP+ (Cao et al., 2025)	w/ QwQ, I	DeepSeek-V3, and BFS-Prover	Informal + - Tree search	671B	1 128 1024	52.5% 74.2% 79.5%
	w/ DeepSe	ek-R1, DeepSeek-V3, and BFS-Prover	- Tice scarcii		1024	80.7%
DeepSeek-Prover-V2 (Re	en et al., 2025	5)	Whole-proof	671B	1 1024	61.9% 86.6%
D 1: D (7)	2025) //	2			8192	88.9%
Delta-Prover (Zhou et al.		Jemini 2.5 Pro	Agent	unknown	16384	95.9%
Seed-Prover (Chen et al.,	2025)		Whole-proof	unknown	unknown	99.6%
Medium Language Mode	ls					
Kimina-Prover-Preview (Whole-proof	72B	1 1024 8192	52.9% 77.9% 80.7%		
Goedel-Prover-V2 (Lin e	t al., 2025c)		Whole-proof	32B	32 1024 8192	88.1% 91.8% 92.2 %
Small Language Models						
DeepSeek-Prover-V1.5-R InternLM2.5-StepProver- HunyuanProver v16 + BF BFS-Prover (Xin et al., 2	Tree search Tree search Tree search Tree search	$\begin{array}{ccc} 7B & 32 \times 16 \times 400 \\ 7B & 256 \times 32 \times 600 \\ 7B & 600 \times 8 \times 400 \\ 7B & 2048 \times 2 \times 600 \end{array}$		63.5% 65.9% 68.4% 70.8%		
Leanabell-Prover-GD-RL Goedel-Prover-SFT (Lin STP (Dong & Ma, 2025)			Whole-proof Whole-proof Whole-proof	7B	128 25600 25600	61.1% 64.7% 67.6%
Kimina-Prover-Preview-Distill (Wang et al., 2025)			Whole-proof		1 32 1024	52.5% 63.1% 70.8%
DeepSeek-Prover-V2 (Re	Whole-proof	7B	1 32 1024 8192	58.6% 75.6% 79.9% 82.0%		
Leanabell-Prover-V2-KM	1 (Ji et al., 20)25)	- Whole-proof	7B	32 128	68.4% 70.4%
Leanabell-Prover-V2-DS	(Ji et al., 202	25)	- whole-proof	/ D	32 128	76.6% 78.2%
Goedel-Prover-V2 (Lin e	Whole-proof	7B	1 64 256 512	60.8% 83.3% 85.2% 85.7%		
w/ DeepSeek	(Direct proving w/o iterative (Direct proving w/o iterative robres) w/ DeepSeek-Prover-V2 (Direct proving w/o iterative robres) (Final proof synthesis w/ lem				1 50 100 260	61.5% 79.9% 82.0% 82.8%
Prover Agent (Ours) w/ Goedel-P	rover-V2	(Direct proving w/o iterative refinement (Direct proving w/o iterative refinement) (Direct proving w/ iterative refinement) (Final proof synthesis w/ lemma)) Agent	8B	1 50 100 260	64.3% 84.4% 85.7% 86.5%
w/ Ensemble Goedel-Pro DeepSeek-	ver-V2 and	(Direct proving w/o iterative refinement (Direct proving w/o iterative refinement) (Direct proving w/ iterative refinement) (Final proof synthesis w/ lemma)			1 50 100 260	64.3% 85.7% 86.9% 88.1 %

approaches that train a single large model end-to-end, our method takes an orthogonal approach by combining an existing LLM and a prover model without any training. This modular design provides a practical benefit, allowing the system to immediately take advantage of improvements in LLMs and prover models by simply replacing components and to scale easily with future advancements.

5.4 EFFECTIVENESS OF INFORMAL, FORMAL, AND LEAN COORDINATION

Table 1 shows that in both model settings, our approach outperforms the corresponding vanilla baselines even before the iterative refinement, highlighting the benefit of collaboration with the informal LLM. Moreover, the scores increase even further after iterative refinement.





(a) Results for different $N_{\rm init}$ and $N_{\rm refine}$. The dotted lines indicate that the corresponding sample budget are used in the proof of a lemma. (b) Histogram of Lean error counts after $N_{\rm init}$.

Figure 3: Ablation study results on $N_{\rm init}$ and $N_{\rm refine}$. These results highlight the importance of initial draft selection and indicate that iterative refinement and lemma-based proving helps overcome saturation from the model's inherent limitations.

5.5 ABLATION STUDY: ANALYZING THE CONTRIBUTION OF EACH STAGE

We conduct an ablation study to illustrate the contribution of each stage of our agent. Results for different $N_{\rm init}$ and $N_{\rm refine}$ are shown in Figure 3a. When $N_{\rm init}$ is set to 1 or 10, the success rate remains significantly lower than that without iterative refinement, even after $N_{\rm refine}=100$ refinement steps. This highlights the importance of the quality of the initial draft used to start refinement: if the initial proof is poor, subsequent corrections become difficult. Comparing $N_{\rm init}=1,10,50$ under the same sample budget shows a clear improvement in performance in this order, indicating the effectiveness of our approach of selecting the proof with the fewest Lean errors. As shown in Figure 3b, the histograms of the minimum number of errors after $N_{\rm init}=1,10,50$ confirm this trend: the error count decreases substantially, and for $N_{\rm init}=50$ most problems have only one or two errors. Although the number of Lean errors may not perfectly measure proof quality, since a single error can still correspond to a mathematically challenging gap, it nevertheless exhibits a strong correlation and serves as a useful proxy for evaluation.

As shown in Figure 3a, the runs without iterative refinement saturate around a sample budget of 80. In contrast, when iterative refinement is applied after $N_{\rm init}=50$ or 100, this saturation is overcome and the success rate improves, outperforming the setting that simply continues generation without refinement. This demonstrates the effectiveness of the iterative refinement: whereas repeated generation alone eventually saturates due to the inherent capacity limits of the model, incorporating external feedback through in-context learning enables the model to improve and overcome this limitation. Also, $N_{\rm init}=50$ and 100 yield almost identical results in the final performance. Since the model had already saturated in this regime, increasing $N_{\rm init}$ did not improve the quality of the selected initial drafts. Furthermore, Figure 3a shows that final synthesis with lemmas improves the score even after iterative refinement has saturated, demonstrating the effectiveness of our lemma-based approach. This indicates that the model's capability is further enhanced by incorporating information beyond mere error feedback.

5.6 CASE STUDY: SUCCESS WITH LEMMA-GUIDED PROOFS

We next present a case study to demonstrate that our approach with auxiliary lemmas is indeed effective in practice. The detailed discussion and the outputs for this problem, such as the generated lemmas, final formal proof, and the associated reasoning process, are provided in Appendix E. We analyze the output and reasoning process for the problem where the direct proof attempt failed but the use of auxiliary lemmas led to a successful proof. In this case, our agent generates a lemma corresponding to the special case of substituting n=3 into the given problem, as well as additional lemmas that may be potentially relevant for solving the problem. As observed in the chain-of-thought process when this lemma is used (see Appendix E.5), the agent immediately considers the n=3 case and then quickly comes up with mathematical induction as the proof strategy. This allows it to quickly transition to filling in the details under a clear proof plan and ultimately complete the proof. Moreover, tactics and proof techniques considered in the auxiliary lemmas reappear in the

Table 2: Comparison of formal theorem-proving performance by problem category on MiniF2F-test. The results are reported as the percentage of theorems proved. The best results in each model setting for each of the three categories, demarcated by double lines, are highlighted in **bold**.

				Olympiad				MATH			Custom			
		Model Size	Sample Budget	IMO	AIME	AMC	Sum	Algebra	Number Theory	Sum	Algebra	Number Theory	Induction	Sum
Number of Problems				20	15	45	80	70	60	130	18	8	8	34
DeepSeek-Prover-V2 (Ren et al., 2025)		671B	8192	50.0	93.3	77.8	73.8	100.0	96.7	98.5	83.3	87.5	100.0	88.2
Prover Agent (Ours) w/ DeepSeek-Prover-V	(Direct proving w/o iterative refinement) (Direct proving w/o iterative refinement) '2 (Direct proving w/ iterative refinement) (Final proof synthesis w/ lemma)	8B	1 50 100 260	40.0 70.0 70.0 70.0	53.3 80.0 80.0 80.0	62.2 82.2 86.7 88.9	55.0 78.8 81.3 82.5	71.4 80.0 84.3 84.3	60.0 88.3 88.3 88.3	66.2 83.8 86.2 86.2	55.6 66.7 66.7 66.7	75.0 75.0 75.0 75.0	50.0 62.5 62.5 75.0	58.8 67.6 67.6 70.6
Goedel-Prover-V2 (Lin et al., 2025c)		8B	1 64 256 512	50.0 80.0 80.0 80.0	60.0 80.0 80.0 80.0	53.3 88.9 88.9 88.9	53.8 85.0 85.0 85.0	71.4 84.3 84.3 84.3	63.3 91.7 91.7 91.7	67.7 87.7 87.7 87.7	50.0 77.8 77.8 77.8	62.5 75.0 75.0 75.0	50.0 87.5 87.5 87. 5	52.9 79.4 79.4 79.4
Prover Agent (Ours) w/ Goedel-Prover-V2	(Direct proving w/o iterative refinement) (Direct proving w/o iterative refinement) (Direct proving w/ iterative refinement) (Final proof synthesis w/ lemma)	8B	1 50 100 260	50.0 80.0 80.0 80.0	73.3 80.0 80.0 80.0	57.8 86.7 88.9 88.9	58.8 83.8 85.0 85.0	68.6 84.3 87.1 88.6	70.0 90.0 90.0 90.0	69.2 86.9 88.5 89.2	55.6 77.8 77.8 77.8	62.5 75.0 75.0 75.0	62.5 75.0 75.0 87.5	58.8 76.5 76.5 79.4
Prover Agent (Ours) w/ Ensenble	(Direct proving w/o iterative refinement) (Direct proving w/o iterative refinement) (Direct proving w/ iterative refinement) (Final proof synthesis w/ lemma)	8B	1 50 100 260	50.0 80.0 80.0 80.0	73.3 80.0 80.0 80.0	57.8 88.9 91.1 93.3	58.8 85.0 86.3 87.5	68.6 87.1 90.0 91.4	70.0 90.0 90.0 90.0	69.2 88.5 90.0 90.8	55.6 77.8 77.8 77.8	62.5 75.0 75.0 75.0	62.5 75.0 75.0 87.5	58.8 76.5 76.5 79.4

reasoning process and final proof: even when a lemma itself is not directly used, the techniques explored during lemma generation provide valuable hints for the overall proof construction. Next, for comparison, we examine the reasoning process without using lemmas, focusing on the trajectory with the fewest final errors (see Appendix E.6). Compared to the successful case with lemmas, the proof strategy here is far less clear, with the model wandering without a coherent plan. As a result, even when it eventually reaches the idea of using mathematical induction, it fails to elaborate the details and the proof does not succeed. This comparison highlights the effectiveness of our auxiliary-lemma approach, which goes beyond simple decomposition of previous work.

5.7 Performance on Olympiad-Level Problems

Table 2 shows the results for each category on the MiniF2F-test dataset. These results demonstrate that our approach with DeepSeek-Prover setting performs particularly well on Olympiad-level problems, even surpassing DeepSeek-Prover-V2 (Ren et al., 2025), which uses a significantly larger 671B model and a much higher sample budget of 8192. Given that our direct proving method without iterative refinement and with a sample budget of only 100 already surpasses DeepSeek-Prover-V2, this suggests that coordination with natural language-based informal reasoning may be the key. Olympiad-level problems require a high degree of mathematical reasoning, and the strong reasoning abilities of the informal LLM likely played a crucial role in solving them effectively. On the other hand, our agent does not outperform DeepSeek-Prover-V2 in the MATH and Custom categories. The consistent gap in these categories suggests that model size and sample budget may play a more significant role here. Since DeepSeek-Prover-V2 also possesses a certain level of mathematical reasoning ability, it can handle these relatively mathematically easier problems on its own. In contrast, with the Goedel-Prover setting, no substantial differences are observed across categories. This is likely because Goedel-Prover already possesses a certain level of the required mathematical capability, and thus category-specific variation does not emerge as clearly.

5.8 Broader Applicability and Future Potential

Nothing in our pipeline is specific to mathematics competition problems. The same approach could be applied to formal proofs in other domains, such as learning theory or physics, as long as the LLM has relevant knowledge or is provided with an appropriate knowledge base. This offers the potential for AI-driven construction of mathematical theories without hallucinations or logical errors.

6 Conclusion

We introduced Prover Agent, a modular framework that coordinates an informal reasoning LLM, a formal prover model, and Lean verification. By generating auxiliary lemmas and leveraging feedback-driven refinement, our method achieved state-of-the-art performance among SLMs on the MiniF2F benchmark.

REFERENCES

- Thomas Anthony, Zheng Tian, and David Barber. Thinking fast and slow with deep learning and tree search. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Zhangir Azerbayev, Bartosz Piotrowski, Hailey Schoelkopf, Edward W. Ayers, Dragomir Radev, and Jeremy Avigad. ProofNet: Autoformalizing and formally proving undergraduate-level mathematics. *arXiv* preprint arXiv:2302.12433, 2023.
- Bruno Barras, Samuel Boutin, Cristina Cornes, Judicaël Courant, Yann Coscoy, David Delahaye, Daniel de Rauglaudre, Jean-Christophe Filliâtre, Eduardo Giménez, Hugo Herbelin, et al. The Coq proof assistant reference manual. *INRIA*, *version*, 6(11):17–21, 1999.
- Chenrui Cao, Liangcheng Song, Zenan Li, Xinyi Le, Xian Zhang, Hui Xue, and Fan Yang. Reviving DSP for advanced theorem proving in the era of reasoning models. *arXiv preprint arXiv:2506.11487*, 2025.
- Luoxin Chen, Jinming Gu, Liankai Huang, Wenhao Huang, Zhicheng Jiang, Allan Jie, Xiaoran Jin, Xing Jin, Chenggang Li, Kaijing Ma, Cheng Ren, Jiawei Shen, Wenlei Shi, Tong Sun, He Sun, Jiahui Wang, Siran Wang, Zhihong Wang, Chenrui Wei, Shufa Wei, Yonghui Wu, Yuchen Wu, Yihang Xia, Huajian Xin, Fan Yang, Huaiyuan Ying, Hongyi Yuan, Zheng Yuan, Tianyang Zhan, Chi Zhang, Yue Zhang, Ge Zhang, Tianyun Zhao, Jianqiu Zhao, Yichi Zhou, and Thomas Hanwen Zhu. Seed-Prover: Deep and broad reasoning for automated theorem proving. *arXiv preprint arXiv:2507.23726*, 2025.
- DeepSeek-AI. DeepSeek-V3 technical report. 2024.
- DeepSeek-AI. DeepSeek-R1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Kefan Dong and Tengyu Ma. STP: Self-play llm theorem provers with iterative conjecturing and proving. *arXiv preprint arXiv:2502.00212*, 2025.
- Kefan Dong, Arvind Mahankali, and Tengyu Ma. Formal theorem proving by rewarding llms to decompose proofs hierarchically. *arXiv preprint arXiv:2411.01829*, 2025.
- Emily First, Markus N. Rabe, Talia Ringer, and Yuriy Brun. Baldur: Whole-proof generation and repair with large language models. ESEC/FSE 2023, pp. 1229–1241. Association for Computing Machinery, 2023. ISBN 9798400703270. doi: 10.1145/3611643.3616243.
- Fabian Gloeckle, Jannis Limperg, Gabriel Synnaeve, and Amaury Hayat. ABEL: Sample efficient online reinforcement learning for neural theorem proving. In *The 4th Workshop on Mathematical Reasoning and AI at NeurIPS*'24, 2024.
- Jesse Michael Han, Jason Rute, Yuhuai Wu, Edward W. Ayers, and Stanislas Polu. Proof artifact co-training for theorem proving with language models. In *International Conference on Learning Representations*, 2022.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. Large language models cannot self-correct reasoning yet. In *The Twelfth International Conference on Learning Representations*, 2024.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Trans. Inf. Syst.*, 43(2), 2025. ISSN 1046-8188. doi: 10.1145/3703155.

- Geoffrey Irving, Christian Szegedy, Alexander A Alemi, Niklas Een, Francois Chollet, and Josef Urban. DeepMath deep sequence models for premise selection. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Xingguang Ji, Yahui Liu, Qi Wang, Jingyuan Zhang, Yang Yue, Rui Shi, Chenxi Sun, Fuzheng Zhang, Guorui Zhou, and Kun Gai. Leanabell-Prover-V2: Verifier-integrated reasoning for formal theorem proving via reinforcement learning. *arXiv preprint arXiv:2507.08649*, 2025.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Comput. Surv.*, 55(12), 2023. ISSN 0360-0300. doi: 10.1145/3571730.
- Albert Q. Jiang, Wenda Li, Jesse Michael Han, and Yuhuai Wu. LISA: Language models of isabelle proofs. In 6th Conference on Artificial Intelligence and Theorem Proving, 2021.
- Albert Q. Jiang, Sean Welleck, Jin Peng Zhou, Wenda Li, Jiacheng Liu, Mateja Jamnik, Timothée Lacroix, Yuhuai Wu, and Guillaume Lample. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. In *The Eleventh International Conference on Learning Represen*tations, 2023.
- Cezary Kaliszyk, Josef Urban, Henryk Michalewski, and Miroslav Olšák. Reinforcement learning of theorem proving. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, pp. 611–626. Association for Computing Machinery, 2023. ISBN 9798400702297. doi: 10.1145/3600006.3613165.
- Guillaume Lample, Timothee Lacroix, Marie-Anne Lachaux, Aurelien Rodriguez, Amaury Hayat, Thibaut Lavril, Gabriel Ebner, and Xavier Martinet. Hypertree proof search for neural theorem proving. In *Advances in Neural Information Processing Systems*, volume 35, pp. 26337–26349. Curran Associates, Inc., 2022.
- Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 3843–3857, 2022.
- Yang Li, Dong Du, Linfeng Song, Chen Li, Weikang Wang, Tao Yang, and Haitao Mi. Hunyuan-Prover: A scalable data synthesis framework and guided tree search for automated theorem proving. *arXiv preprint arXiv:2412.20735*, 2025.
- Haohan Lin, Zhiqing Sun, Sean Welleck, and Yiming Yang. Lean-STaR: Learning to interleave thinking and proving. In *The Thirteenth International Conference on Learning Representations*, 2025a.
- Yong Lin, Shange Tang, Bohan Lyu, Jiayun Wu, Hongzhou Lin, Kaiyu Yang, Jia Li, Mengzhou Xia, Danqi Chen, Sanjeev Arora, and Chi Jin. Goedel-Prover: A frontier model for open-source automated theorem proving. *arXiv preprint arXiv:2502.07640*, 2025b.
- Yong Lin, Shange Tang, Bohan Lyu, Ziran Yang, Jui-Hui Chung, Haoyu Zhao, Lai Jiang, Yihan Geng, Jiawei Ge, Jingruo Sun, Jiayun Wu, Jiri Gesi, Ximing Lu, David Acuna, Kaiyu Yang, Hongzhou Lin, Yejin Choi, Danqi Chen, Sanjeev Arora, and Chi Jin. Goedel-Prover-V2: Scaling formal theorem proving with scaffolded data synthesis and self-correction. *arXiv preprint arXiv:2508.03613*, 2025c.
- Pan Lu, Liang Qiu, Wenhao Yu, Sean Welleck, and Kai-Wei Chang. A survey of deep learning for mathematical reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14605–14631. Association for Computational Linguistics, July 2023. doi: 10.18653/v1/2023.acl-long.817.

- The mathlib Community. The lean mathematical library. In *Proceedings of the 9th ACM SIG-PLAN International Conference on Certified Programs and Proofs*, pp. 367–381. Association for Computing Machinery, 2020. ISBN 9781450370974. doi: 10.1145/3372885.3373824.
 - Norman D. Megill and David A. Wheeler. *Metamath: A Computer Language for Pure Mathematics*, 2019. URL http://us.metamath.org/downloads/metamath.pdf.
 - Leonardo de Moura and Sebastian Ullrich. 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*, LNCS 12699, pp. 625–635. Springer-Verlag, 2021. doi: 10.1007/978-3-030-79876-5_37.
 - A. Newell and H. Simon. The logic theory machine–a complex information processing system. *IRE Transactions on Information Theory*, 2(3):61–79, 1956. doi: 10.1109/TIT.1956.1056797.
 - OpenAI. OpenAI o1 system card. arXiv preprint arXiv:2412.16720, 2024.
 - OpenAI. OpenAI o3-mini, January 2025. URL https://openai.com/index/openai-o3-mini/.
 - Lawrence C. Paulson. Isabelle a Generic Theorem Prover. Springer Verlag, 1994.
 - Stanislas Polu and Ilya Sutskever. Generative language modeling for automated theorem proving. *arXiv preprint arXiv:2009.03393*, 2020.
 - Stanislas Polu, Jesse Michael Han, Kunhao Zheng, Mantas Baksys, Igor Babuschkin, and Ilya Sutskever. Formal mathematics statement curriculum learning. In *The Eleventh International Conference on Learning Representations*, 2023.
 - Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems*, volume 36, pp. 53728–53741, 2023.
 - Z. Z. Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanjia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, Z. F. Wu, Zhibin Gou, Shirong Ma, Hongxuan Tang, Yuxuan Liu, Wenjun Gao, Daya Guo, and Chong Ruan. DeepSeek-Prover-V2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. 2025.
 - Ziju Shen, Naohao Huang, Fanyi Yang, Yutong Wang, Guoxiong Gao, Tianyi Xu, Jiedong Jiang, Wanyi He, Pu Yang, Mengzhou Sun, Haocheng Ju, Peihao Wu, Bryan Dai, and Bin Dong. REAL-Prover: Retrieval augmented Lean prover for mathematical reasoning. *arXiv preprint arXiv:2505.20613*, 2025.
 - Yuda Song, Hanlin Zhang, Carson Eisenach, Sham M. Kakade, Dean Foster, and Udaya Ghai. Mind the gap: Examining the self-improvement capabilities of large language models. In *The Thirteenth International Conference on Learning Representations*, 2025.
 - Kaya Stechly, Karthik Valmeekam, and Subbarao Kambhampati. On the self-verification limitations of large language models on reasoning and planning tasks. In *The Thirteenth International Conference on Learning Representations*, 2025.
 - Amitayush Thakur, George Tsoukalas, Yeming Wen, Jimmy Xin, and Swarat Chaudhuri. An incontext learning agent for formal theorem-proving. In *First Conference on Language Modeling*, 2024.
 - George Tsoukalas, Jasper Lee, John Jennings, Jimmy Xin, Michelle Ding, Michael Jennings, Amitayush Thakur, and Swarat Chaudhuri. Putnambench: Evaluating neural theorem-provers on the putnam mathematical competition. In *Advances in Neural Information Processing Systems*, volume 37, pp. 11545–11569, 2024.

Haiming Wang, Ye Yuan, Zhengying Liu, Jianhao Shen, Yichun Yin, Jing Xiong, Enze Xie, Han Shi, Yujun Li, Lin Li, Jian Yin, Zhenguo Li, and Xiaodan Liang. DT-solver: Automated theorem proving with dynamic-tree sampling guided by proof-level value function. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12632–12646. Association for Computational Linguistics, 2023. doi: 10.18653/v1/2023. acl-long.706.

- Haiming Wang, Huajian Xin, Zhengying Liu, Wenda Li, Yinya Huang, Jianqiao Lu, Zhicheng Yang, Jing Tang, Jian Yin, Zhenguo Li, and Xiaodan Liang. Proving theorems recursively. In *Advances in Neural Information Processing Systems*, volume 37, pp. 86720–86748, 2024a.
- Haiming Wang, Mert Unsal, Xiaohan Lin, Mantas Baksys, Junqi Liu, Marco Dos Santos, Flood Sung, Marina Vinyes, Zhenzhe Ying, Zekai Zhu, Jianqiao Lu, Hugues de Saxcé, Bolton Bailey, Chendong Song, Chenjun Xiao, Dehao Zhang, Ebony Zhang, Frederick Pu, Han Zhu, Jiawei Liu, Jonas Bayer, Julien Michel, Longhui Yu, Léo Dreyfus-Schmidt, Lewis Tunstall, Luigi Pagani, Moreira Machado, Pauline Bourigault, Ran Wang, Stanislas Polu, Thibaut Barroyer, Wen-Ding Li, Yazhe Niu, Yann Fleureau, Yangyang Hu, Zhouliang Yu, Zihan Wang, Zhilin Yang, Zhengying Liu, and Jia Li. Kimina-prover preview: Towards large formal reasoning models with reinforcement learning. *arXiv preprint arXiv:2504.11354*, 2025.
- Ruida Wang, Jipeng Zhang, Yizhen Jia, Rui Pan, Shizhe Diao, Renjie Pi, and Tong Zhang. TheoremLlama: Transforming general-purpose LLMs into lean4 experts. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 11953–11974. Association for Computational Linguistics, 2024b. doi: 10.18653/v1/2024.emnlp-main.667.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 24824–24837, 2022.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45. Association for Computational Linguistics, 2020.
- Minchao Wu, Michael Norrish, Christian Walder, and Amir Dezfouli. Tacticzero: Learning to prove theorems from scratch with deep reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 34, pp. 9330–9342, 2021.
- Zijian Wu, Suozhi Huang, Zhejian Zhou, Huaiyuan Ying, Jiayu Wang, Dahua Lin, and Kai Chen. InternLM2.5-StepProver: Advancing automated theorem proving via expert iteration on large-scale lean problems. 2024a.
- Zijian Wu, Jiayu Wang, Dahua Lin, and Kai Chen. LEAN-GitHub: Compiling github lean repositories for a versatile lean prover. *arXiv preprint arXiv:2407.17227*, 2024b.
- Huajian Xin, Z. Z. Ren, Junxiao Song, Zhihong Shao, Wanjia Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, Wenjun Gao, Qihao Zhu, Dejian Yang, Zhibin Gou, Z. F. Wu, Fuli Luo, and Chong Ruan. DeepSeek-Prover-V1.5: Harnessing proof assistant feedback for reinforcement learning and monte-carlo tree search, 2025a.
- Ran Xin, Chenguang Xi, Jie Yang, Feng Chen, Hang Wu, Xia Xiao, Yifan Sun, Shen Zheng, and Kai Shen. BFS-Prover: Scalable best-first tree search for llm-based automatic theorem proving. 2025b.
- Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*, 2025.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. Qwen2.5-Math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024.

- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report. 2025a.
- Kaiyu Yang, Aidan Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil, Ryan J Prenger, and Animashree Anandkumar. LeanDojo: Theorem proving with retrieval-augmented language models. In *Advances in Neural Information Processing Systems*, volume 36, pp. 21573–21612, 2023.
- Kaiyu Yang, Gabriel Poesia, Jingxuan He, Wenda Li, Kristin E. Lauter, Swarat Chaudhuri, and Dawn Song. Position: Formal mathematical reasoning—a new frontier in AI. In *Forty-second International Conference on Machine Learning Position Paper Track*, 2025b.
- Huaiyuan Ying, Zijian Wu, Yihan Geng, Jiayu Wang, Dahua Lin, and Kai Chen. Lean workbook: A large-scale lean problem set formalized from natural language math problems. In *Advances in Neural Information Processing Systems*, volume 37, pp. 105848–105863, 2024.
- Roozbeh Yousefzadeh and Xuenan Cao. A lean dataset for international math olympiad: Small steps towards writing math proofs for hard problems. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856.
- Zhiyuan Zeng, Qinyuan Cheng, Zhangyue Yin, Yunhua Zhou, and Xipeng Qiu. Revisiting the test-time scaling of o1-like models: Do they truly possess test-time scaling capabilities? 2025.
- Jingyuan Zhang, Qi Wang, Xingguang Ji, Yahui Liu, Yang Yue, Fuzheng Zhang, Di Zhang, Guorui Zhou, and Kun Gai. Leanabell-Prover: Posttraining scaling in formal reasoning. *arXiv* preprint *arXiv*:2504.06122, 2025.
- Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. miniF2F: a cross-system benchmark for formal olympiad-level mathematics. In *International Conference on Learning Representations*, 2022.
- Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Language agent tree search unifies reasoning, acting, and planning in language models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pp. 62138–62160. PMLR, 2024.
- Yichi Zhou, Jianqiu Zhao, Yongxin Zhang, Bohan Wang, Siran Wang, Luoxin Chen, Jiahui Wang, Haowei Chen, Allan Jie, Xinbo Zhang, Haocheng Wang, Luong Trung, Rong Ye, Phan Nhat Hoang, Huishuai Zhang, Peng Sun, and Hang Li. Solving formal math problems by decomposition and iterative reflection. *arXiv preprint arXiv:2507.15225*, 2025.

A EXTENDED RELATED WORK

 We briefly summarized related work in Section 2. Here we provide details of representative systems.

A.1 LANGUAGE MODELS FOR FORMAL THEOREM PROVING

The use of language models for guiding formal theorem provers has gained momentum recently. Early work like GPT-f (Polu & Sutskever, 2020) applied transformers to produce proofs in formal systems, such as Metamath (Megill & Wheeler, 2019) and Lean (Moura & Ullrich, 2021), by generating one proof step (tactic) at a time, guided by a goal state. Subsequent efforts in Lean, such as lean-gptf⁴ and PACT (Han et al., 2022), fine-tuned LLMs on large corpora of proof data, achieving moderate success in automatically discovering proofs.

A.2 TREE-SEARCH-BASED FORMAL PROVING

BFS-Prover (Xin et al., 2025b) proposed a scalable best-first tree search framework for Lean 4 that incorporates three key innovations: strategic data filtering during expert iterations, direct preference optimization (DPO) (Rafailov et al., 2023) on state-tactic pairs using Lean compiler feedback, and length normalization to encourage exploration of deeper proof paths. InternLM2.5-StepProver (Wu et al., 2024a) combined expert iteration with BFS and critic-guided sampling, while Hunyuan-Prover (Li et al., 2025) integrated large-scale data synthesis and guided search. Reinforcement-enhanced variants such as DeepSeek-Prover-V1.5 (Xin et al., 2025a) proposed the use of RMaxTS, a variant of Monte-Carlo tree search (MCTS), to diversify exploration and improve success rates.

A.3 WHOLE-PROOF GENERATION

Representative systems in this strand have advanced two complementary mechanisms: (i) expertiteration bootstrapping, which cycles model-generated proofs through a formal verifier to curate training trajectories, and (ii) reinforcement learning (RL) with verifier feedback that directly optimizes long, one-shot scripts (often with a long chain-of-thought).

Polu et al. (2023) introduced expert iteration for formal mathematics, alternating proof search with learning. They showed expert iteration outperforms search-only at fixed compute, discovered an automatically paced curriculum from problem statements, and showed improved performance on the miniF2F (Zheng et al., 2022) benchmark without requiring ground-truth proofs. InternLM2.5-StepProver (Wu et al., 2024a) scaled expert iteration on Lean-Workbook (Ying et al., 2024), trained a critic to prioritize easier instances and guide deeper proofs, and paired expert iteration with best-first exploration, achieving strong results on several benchmarks, such as miniF2F (Zheng et al., 2022), ProofNet (Azerbayev et al., 2023), PutnamBench (Tsoukalas et al., 2024), and Lean-Workbook-Plus (Ying et al., 2024). Lean-STaR (Lin et al., 2025a) trained a model to interleave informal natural-language thoughts with formal tactic steps. The model is trained by expert iteration, and at inference time, it generates informal reasoning prior to each tactic, enhancing theorem-proving performance. Goedel-Prover (Lin et al., 2025b) tackled data scarcity by training statement formalizers to translate Numina problems into Lean 4, building a 1.64M-statement corpus, and iteratively bootstrapping provers whose new proofs are added to training. The resulting SFT-centered expert iteration pipeline surpasses prior open-source baselines. Goedel-Prover-V2 (Lin et al., 2025c) extended expert iteration with scaffolded data synthesis, verifier-guided self-correction, and model averaging, delivering large gains on the MiniF2F benchmark (Zheng et al., 2022) at 8-32B scales under constrained test-time budgets.

Kaliszyk et al. (2018) formulated theorem proving as reinforcement learning for connection-style proof search, using Monte Carlo simulations guided by rewards from previous attempts to replace hand-crafted heuristics and improve held-out performance. DeepSeek-Prover-V1.5 (Xin et al., 2025a) utilized reinforcement learning from proof assistant feedback (RLPAF) and a novel Monte-Carlo tree search variant, RMaxTS, which employs an intrinsic-reward-driven strategy to explore diverse proof paths. Leanabell-Prover (Zhang et al., 2025) demonstrated the effectiveness of post-training in formal theorem proving by applying continual training with data emulating human cog-

⁴https://github.com/jesse-michael-han/lean-gptf

nitive behaviors and reinforcement learning with compiler feedback to existing models. Kimina-Prover Preview (Wang et al., 2025) employed a large-scale reinforcement learning pipeline and a structured "formal reasoning pattern," emulating human problem-solving strategies. It achieves an 80.7% pass rate on MiniF2F (Zheng et al., 2022) with a 72B-parameter model. Leanabell-Prover-V2 (Ji et al., 2025) is built on Kimina-Prover-Preview-Distill-7B(Wang et al., 2025) and DeepSeek-Prover-V2-7B (Ren et al., 2025) as base models, and further improved through post-training with reinforcement learning.

A.4 FORMAL THEOREM PROVING WITH RETRIEVAL-AUGMENTED GENERATION

Retrieval-augmented provers query large formal libraries at inference time and condition generation on the retrieved items, typically relevant lemmas, theorems, or proof patterns from mathlib (mathlib Community, 2020). This mitigates the limits of parametric memory by injecting on-demand knowledge and can be applied to both stepwise tactic generation and whole-proof scripts. Lean-Dojo (Yang et al., 2023) established the core infrastructure for RAG in Lean, including fine-grained premise annotations, a gym-like interactive environment, and a retrieval-augmented prover that selects premises for each proof state. REAL-Prover (Shen et al., 2025) integrated a semantic premise selector (LeanSearch-PS) with a fine-tuned Lean 4 prover and reports gains on challenging benchmarks such as ProofNet (Azerbayev et al., 2023).

A.5 PROOF REFINEMENT AND SUBGOAL DECOMPOSITION

Jiang et al. (2023) introduced Draft, Sketch, and Prove (DSP), a novel three-stage method that leverages informal proofs to guide automated theorem provers. The process involves drafting an informal proof (either by a human or an LLM), using a language model to convert it into a high-level formal sketch with verifiable steps, and finally employing an off-the-shelf prover to automatically solve the remaining logical gaps. This approach of guiding a formal prover with an informal-to-formal sketch significantly improved its success rate, boosting performance on the miniF2F benchmark from 20.9% to 39.3%.

Wang et al. (2024a) introduced POETRY, a novel method that proves theorems recursively to overcome the limitations of short-sighted, step-by-step search in automated theorem proving. By first finding a verifiable high-level proof sketch and deferring detailed sub-proofs to subsequent levels using a *sorry* tactic, POETRY can solve more complex problems and find significantly longer proofs, leading to superior results on the miniF2F (Zheng et al., 2022) and PISA (Jiang et al., 2021) benchmarks.

Cao et al. (2025) introduced DSP+, an improved Draft, Sketch, and Prove framework Jiang et al. (2023) that achieves high performance in automated theorem proving without requiring any model training or fine-tuning. By carefully coordinating existing off-the-shelf reasoning models and step provers with fine-grained neuro-symbolic enhancements at each stage, DSP+ solved 80.7% of the miniF2F benchmark (Zheng et al., 2022), which was comparable to top models that rely on extensive reinforcement learning, and even proved a previously unsolved IMO problem.

DeepSeek-Prover-V2 (Ren et al., 2025) used a powerful general-purpose model, DeepSeek-V3 (DeepSeek-AI, 2024), to break down complex theorems into simpler subgoals, which are then recursively solved and synthesized into a cold-start dataset for the final prover. The resulting model achieved an 88.9% pass rate on the MiniF2F benchmark (Zheng et al., 2022).

Delta Prover (Zhou et al., 2025) is an agent-based framework that enables a general-purpose LLM to solve formal math problems without any specialized fine-tuning. The agent orchestrated the LLM's interaction with the Lean 4 environment through a novel process of reflective decomposition and iterative proof repair, where the model breaks down complex problems and corrects its own errors based on compiler feedback. This training-free approach achieved a 95.9% success rate on the miniF2F benchmark (Zheng et al., 2022), surpassing all previous methods, including those requiring extensive specialized training.

Chen et al. (2025) introduced Seed-Prover, a whole-proof reasoning model that uses a novel lemmastyle approach to solve complex formal math problems. Seed-Prover iteratively refined its proofs using compiler feedback and a shared pool of proved lemmas, employing a powerful three-tiered test-time inference strategy for both deep and broad reasoning. This method significantly surpassed

866

867

868

870

871

872

873

874

875

876

877

878

879

880

882

883

885

886

887

889

890

891

892

893

894

895

896

897 898 899

900

901 902 903

904 905

906907908

909 910

911

912 913

914 915

916

917

Algorithm 1 The overall architecture of our lemma-based theorem-proving agent coordinating informal reasoning, formal reasoning, and Lean.

```
Input: Problem T with hyperparameters N_{\text{init}} (max initial proof attempts) and N_{\text{refine}} (max refinement attempts)
Output: Formal proof of T or failure
function MAIN(T): Overall proof process for problem T
                                                                    function PROVE(S): Attempt to generate an informal
    P_{\text{direct.}} \leftarrow PROVE(T): Attempt to prove theorem T
                                                                    proof of S
directly
                                                                        // Initial proof attempt
   if P_{\text{direct}} succeeds then
                                                                        for k = 1 to N_{\text{init}} do
        return P_{
m direct}
                                                                            Informal LLM generates informal proof P_{\inf} of S
    end if
                                                                            Prover attempts to formalize P_{inf} into P_{form}
                                                                            Lean checks P_{\mathrm{form}}
    // Generate lemmas
    Informal LLM generates lemmas L_1, L_2, \ldots, L_n in
                                                                            if the check succeeds then
                                                                                return P_{
m form}
natural language
    for each lemma L_i do
                                                                            end if
        AutoFormalizer converts L_i into Lean statement
                                                                        end for
F_i
                                                                        // Iterative refinement
        Lean checks F_i. If failing, regenerate F_i until syn-
                                                                        P_{\text{best}} \leftarrow \text{Best previous proof attempt with the fewest}
tactically correct
                                                                    Lean errors
                                                                        return IterativeRefine(P_{best})
    end for
   // Prove each lemma
                                                                    end function
   for each lemma F_i do
        P_i \leftarrow Prove(F_i): Attempt to prove lemma F_i
                                                                    function ITERATIVEREFINE(P): Refine proof P based on
    end for
                                                                    Lean feedback
   // Collect proven lemmas
                                                                        for k=1 to N_{\text{refine}} do
                                                                            Prover generates revised proof P' based on Lean
    \mathcal{P}_{\text{proven}} \leftarrow \{P_i \mid P_i \text{ is succeeded}\}
   // Synthesize final proof using proven lemmas
                                                                    feedback
    for k=1 to N_{\rm init} do
                                                                            Lean checks P'
        P_{\text{final}} \leftarrow \text{Prover synthesizes proof of } T \text{ using}
                                                                            if the check succeeds then
                                                                                return P'
        Lean checks P_{\text{final}}
                                                                                P \leftarrow P' // Update best proof
        if the check succeeds then
            return P_{\rm final}
                                                                            end if
        end if
                                                                        end for
    end for
                                                                        return failure // No proof found after max attempts
   // Iterative refinement of final proof
                                                                    end function
    P_{\text{best}} \leftarrow \text{Best previous proof attempt with the fewest}
    return IterativeRefine(P_{\text{best}})
end function
```

all previous state-of-the-art results, saturating the MiniF2F benchmark (Zheng et al., 2022), proving 78.1% of past IMO problems, and solving 5 out of 6 problems at the IMO 2025 competition.

B PSEUDOCODE OF THE OVERALL WORKFLOW

The pseudocode of our overall workflow is shown in Algorithm 1.

C DETAILED THEORETICAL ANALYSIS

We briefly discussed the theoretical analysis of our approach in Section 4. In this section, we provide a detailed theoretical analysis of our approach.

C.1 Benefits of Lemmas for Structured Proof Decomposition

We begin by stating a lemma required for the following analysis:

Lemma C.1 (Number of Trials for Success). Let p denote the probability that the model successfully proves a theorem T. Then the expected number of trials until the first success, N, and the number

of trials required to succeed with probability at least $1 - \delta$, denoted N_{δ} , satisfy the following:

$$\mathbb{E}[N] = \frac{1}{p}, \quad \log(1/\delta) \left(\frac{1}{p} - 1\right) < \frac{\log \delta}{\log(1-p)} < N_{\delta} = \left\lceil \frac{\log \delta}{\log(1-p)} \right\rceil < \frac{\log(1/\delta)}{p} + 1.$$

Proof. Since each trial is an independent Bernoulli experiment with success probability p, the number of trials N until the first success follows a geometric distribution. It is well known that

$$\mathbb{E}[N] = \sum_{n=1}^{\infty} n(1-p)^{n-1}p = \frac{1}{p}.$$

Next, we consider N_{δ} . Since the probability of at least one success in n trials is $1 - (1 - p)^n$, the condition for achieving success with probability at least $1 - \delta$ is:

$$1 - (1 - p)^n = 1 - \delta \iff (1 - p)^n = \delta \iff n = \frac{\log \delta}{\log(1 - p)}.$$

Recalling the standard inequalities $p \le -\log(1-p) \le \frac{p}{1-p}$, which is valid for $0 , together with the basic ceiling inequality <math>x \le \lceil x \rceil < x+1$, we obtain:

$$\log(1/\delta)\left(\frac{1}{p}-1\right) < \frac{\log \delta}{\log(1-p)} < N_{\delta} = \left\lceil \frac{\log \delta}{\log(1-p)} \right\rceil < \frac{\log(1/\delta)}{p} + 1.$$

This completes the proof.

For simplicity, we henceforth relax N_{δ} to be continuous and write:

$$\log(1/\delta)\left(\frac{1}{p} - 1\right) < N_{\delta} = \frac{\log \delta}{\log(1-p)} < \frac{\log(1/\delta)}{p}.$$

The difference from the actual integer-valued N_{δ} is at most less than 1.

As rigorous versions of Theorems 4.4 to 4.5 described in Section 4.1, we obtain the following Theorems C.2 to C.3, under the same Assumptions 4.1 to 4.3:

Theorem C.2 (Required Number of Trials). Let N_{dir} denote the number of trials required to directly prove a problem T with probability at least $1-\delta$. Let N_{lem} denote the total number of trials required to complete the proof of T with probability at least $1-\delta$, when lemmas L_1,\ldots,L_n are introduced with an allowed failure probability δ_{lem} . Suppose each lemma L_i contains a subset of the essential intermediate facts $\{F_i\}_{i\in S_i}$ with $S_i\subseteq [m]$. Then the following holds:

$$\Phi_{\rm dir}(p) - \log(1/\delta) < N_{\rm dir} < \Phi_{\rm dir}(p),$$

$$\Phi_{\rm lem}(p) - \log(1/\delta) - n\log(1/\delta_{\rm lem}) < \mathbb{E}[N_{\rm lem}] < \Phi_{\rm lem}(p),$$

where

$$\Phi_{\mathrm{dir}}(p) \coloneqq \log(1/\delta) \prod_{i=1}^{m} \frac{1}{p_i},$$

$$\Phi_{\text{lem}}(p) \coloneqq \log(1/\delta_{\text{lem}}) \sum_{i=1}^{n} \prod_{j \in S_i} \frac{1}{p_j} + \frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \prod_{i=1}^{n} \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \prod_{j \in S_i} \frac{1}{p_j} \right).$$

Here, we denote $U := \bigcup_{i=1}^n S_i$, $R_0 := [m] \setminus U$, and $r_0 := \min P(F_S | \{F_i\}_{i \in S})$.

Proof. By Assumption 4.2, the probability that all F_1, \ldots, F_m succeed and the problem T is solved equals $\prod_{i=1}^m p_i$. Hence, by Lemma C.1, we obtain:

$$\Phi_{\rm dir}(p) - \log(1/\delta) < N_{\rm dir} < \Phi_{\rm dir}(p).$$

Similarly, since the probability that all F_j with $j \in S_i$ succeed and lemma L_i is proved equals $\prod_{j \in S_i} p_j$, the number of trials required for lemma L_i , denoted N_{L_i} , satisfies:

$$\log(1/\delta_{\text{lem}}) \prod_{j \in S_i} \frac{1}{p_j} - \log(1/\delta_{\text{lem}}) < N_{L_i} < \log(1/\delta_{\text{lem}}) \prod_{j \in S_i} \frac{1}{p_j}.$$

Therefore, the total number of trials required to prove all n lemmas L_1, \ldots, L_n is bounded by the sum of the bounds above, i.e.,

$$\log(1/\delta_{\text{lem}}) \sum_{i=1}^{n} \prod_{j \in S_i} \frac{1}{p_j} - n \log(1/\delta_{\text{lem}}) < \sum_{i=1}^{n} N_{L_i} < \log(1/\delta_{\text{lem}}) \sum_{i=1}^{n} \prod_{j \in S_i} \frac{1}{p_j}.$$
 (1)

The probability that the composition of all lemmas succeeds is r_0 , while the probability of proving the uncovered facts $\{F_i\}_{i\in R_0}$ is $\prod_{i\in R_0}p_i$. If a lemma L_i fails with probability δ_{lem} , then in the final proof it must be reproved directly, which succeeds with probability $\prod_{j\in S_i}p_j$. Thus, the expected success probability of lemma L_i in the final stage is: $(1-\delta_{\mathrm{lem}})+\delta_{\mathrm{lem}}\prod_{i\in S_i}p_i$.

Therefore, since the expected success probability in the final stage is given by the product above, the number of trials required to complete the proof of the whole problem T using lemmas in the final stage, denoted $N_{\rm final}$, satisfies:

$$\Phi_{\text{final}}(p) - \log(1/\delta) < \mathbb{E}[N_{\text{final}}] < \Phi_{\text{final}}(p), \tag{2}$$

where

$$\Phi_{\mathrm{final}}(p) \coloneqq \frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \prod_{i=1}^n \left((1 - \delta_{\mathrm{lem}}) + \delta_{\mathrm{lem}} \prod_{j \in S_i} \frac{1}{p_j} \right).$$

Hence, by combining Equations (1) and (2), we obtain the desired result, completing the proof of Theorem C.2. \Box

From Theorem C.2, we see that decomposing the problem into lemmas transforms the corresponding leading term from a product into a sum, thereby significantly reducing the order of the required number of trials.

Theorem C.3 (Threshold Condition for Lemma Efficiency). There exists a threshold $\tau \in [0,1]$ such that if $p_i \leq \tau$ for all $i \in [m]$, then $\mathbb{E}[N_{\text{lem}}] \leq N_{\text{dir}}$ holds for any $\delta, \delta_{\text{lem}} \in (0,1)$.

Proof. Consider the condition $\frac{\mathbb{E}[N_{\text{lem}}]}{N_{\text{dir}}} < 1$. By Theorem C.2, this condition is satisfied if the following holds:

$$\frac{\Phi_{\text{lem}}(p)}{\Phi_{\text{dir}}(p) - \log(1/\delta)} < 1$$

$$\Leftrightarrow \frac{\log(1/\delta_{\text{lem}}) \sum_{i=1}^{n} \prod_{j \in S_i} \frac{1}{p_j}}{\log(1/\delta) \prod_{i=1}^{m} \left(\frac{1}{p_i} - 1\right)}$$

$$+ \frac{\frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i}\right) \prod_{i=1}^{n} \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \prod_{j \in S_i} \frac{1}{p_j}\right)}{\log(1/\delta) \prod_{i=1}^{m} \left(\frac{1}{p_i} - 1\right)} < 1.$$
(3)

The first term on the left-hand side (LHS) of Equation (3) can be rewritten as:

$$\frac{\log(1/\delta_{\text{lem}}) \sum_{i=1}^{n} \prod_{j \in S_{i}} \frac{1}{p_{j}}}{\log(1/\delta) \prod_{i=1}^{m} \left(\frac{1}{p_{i}} - 1\right)} = \frac{\log(1/\delta_{\text{lem}})}{\log(1/\delta)} \sum_{i=1}^{n} \frac{\prod_{j \in S_{i}} \frac{1}{p_{j}} \prod_{j=1}^{m} p_{j}}{1 - \prod_{j=1}^{m} p_{j}}$$

$$= \frac{\log(1/\delta_{\text{lem}})}{\log(1/\delta)} \sum_{i=1}^{n} \frac{\prod_{j \notin S_{i}} p_{j}}{1 - \prod_{j=1}^{m} p_{j}}.$$
(4)

The second term on the LHS of Equation (3) can be rewritten as:

$$\frac{\frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \prod_{i=1}^n \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \prod_{j \in S_i} \frac{1}{p_j} \right)}{\log(1/\delta) \prod_{i=1}^m \left(\frac{1}{p_i} - 1 \right)}$$

$$= \frac{1}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \left(\prod_{i=1}^n \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \prod_{j \in S_i} \frac{1}{p_j} \right) \right) \frac{\prod_{j=1}^m p_j}{1 - \prod_{j=1}^m p_j}$$

$$= \frac{1}{r_0} \prod_{i=1}^n \left((1 - \delta_{\text{lem}}) \prod_{j \in S_i} p_j + \delta_{\text{lem}} \right) \frac{1}{1 - \prod_{j=1}^m p_j}.$$
(5)

From Equations (4) and (5), both the first and second terms on the LHS of Equation (3) are monotonically increasing with respect to p_i . Hence, the LHS of Equation (3) itself is monotonically increasing w.r.t. p_i . Therefore, by bounding the LHS of Equation (3) from above by using $p_{\max} := \max_i p_i$ and solving for p_{\max} , we obtain a sufficient condition, completing the proof.

From Theorem C.3, it follows that lemma generation is effective for difficult problems. Therefore, our strategy of generating lemmas for difficult problems and solving easy problems directly is justified.

Theorem C.4 (Optimal Partition of Lemma Coverage). Under the fixed lemma coverage $U := \bigcup_{i=1}^n S_i \subseteq [m]$, $\mathbb{E}[N_{\text{lem}}]$ is minimized when $\log p(S_i)$ is as close as possible to $\frac{1}{n} \log p(U)$ for all $i \in [n]$, where $p(S_i) := \prod_{j \in S_i} p_j$ and $p(U) := \prod_{j \in U} p_j$.

Proof. From Theorem C.2, we consider minimizing $\Phi_{lem}(p)$. Let $W := \prod_{i \in U} \frac{1}{p_i}$.

By Jensen's inequality, the first term of $\Phi_{lem}(p)$ can be bounded as follows:

$$\log(1/\delta_{\text{lem}}) \sum_{i=1}^{n} \prod_{j \in S_{i}} \frac{1}{p_{j}} = \log(1/\delta_{\text{lem}}) \sum_{i=1}^{n} \exp(\sum_{j \in S_{i}} \log \frac{1}{p_{j}})$$

$$\geq \log(1/\delta_{\text{lem}}) n \exp\left(\frac{1}{n} \sum_{i=1}^{n} \sum_{j \in S_{i}} \log \frac{1}{p_{j}}\right)$$

$$= \log(1/\delta_{\text{lem}}) n \exp\left(\frac{1}{n} \log W\right)$$

with equality if and only if $\log p(S_i) = \frac{1}{n} \log p(U)$ for all $i \in [n]$.

Noting that $f(x) = \log((1-d) + d \exp(x))$ is convex for $d \in (0,1)$, we can apply Jensen's inequality to bound the second term of $\Phi_{lem}(p)$ as follows:

$$\frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \prod_{i=1}^n \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \prod_{j \in S_i} \frac{1}{p_j} \right) \\
= \frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \exp\left(\sum_{i=1}^n \log \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \exp\left(\sum_{j \in S_i} \log \frac{1}{p_j} \right) \right) \right) \\
\ge \frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \exp\left(n \log \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \exp\left(\frac{1}{n} \sum_{i=1}^n \sum_{j \in S_i} \log \frac{1}{p_j} \right) \right) \right) \\
= \frac{\log(1/\delta)}{r_0} \left(\prod_{i \in R_0} \frac{1}{p_i} \right) \exp\left(n \log \left((1 - \delta_{\text{lem}}) + \delta_{\text{lem}} \exp\left(\frac{1}{n} \log W \right) \right) \right)$$

with equality if and only if $\log p(S_i) = \frac{1}{n} \log p(U)$ for all $i \in [n]$.

Therefore, since both the first and second terms of $\Phi_{lem}(p)$ attain their minimum under the same condition, namely:

 $\log p(S_i) = \frac{1}{n} \log p(U)$ for all $i \in [n]$,

it follows that $\Phi_{\rm lem}(p)$ itself is minimized under this condition. In the discrete case, the minimum is achieved at the partition closest to this balanced condition. This completes the proof.

Theorem C.4 suggests that the optimal lemmas are those that divide the problem into subproblems of approximately equal difficulty.

C.2 Benefits of Lemmas for Discovering Proof Strategies (e.g., Specific Cases)

Theorem C.5 (Success Probability Improvement by Lemmas (Restated)). The success probability of performing one trial of final proving by sampling a strategy from the posterior distribution π_n is bounded as follows:

$$\mathbb{E}[\mathbb{P}(\mathsf{succ}@1)] \geq r \exp(-H_0 + I(Z; Y_{1:n})).$$

Proof. We begin with:

$$\mathbb{P}(\text{succ}@1 \mid Z = z, Y = y) = p(z) \pi(z \mid y).$$

Taking expectation, we obtain:

$$\begin{split} \mathbb{E}_{Z,Y} \left[\mathbb{P}(\mathsf{succ}@1 \mid Z, Y) \right] &= \mathbb{E}_{Z,Y} \left[p(Z) \, \pi(Z \mid Y) \right] \\ &= \mathbb{E}_{Z,Y} \left[p(Z) \, \pi_n(Z) \right] \\ &\geq r \, \mathbb{E}_{Z,Y} [\pi_n(Z)]. \end{split} \tag{6}$$

It remains to lower-bound $\mathbb{E}_{Z,Y}[\pi_n(Z)]$.

For fixed Y = y, we have:

$$\mathbb{E}_{Z} \left[\pi_{n}(Z) \mid Y = y \right] = \sum_{z \in \mathcal{S}} \pi_{n}(z) \, \mathbb{P}(Z = z \mid Y = y)$$
$$= \sum_{z \in \mathcal{S}} \pi_{n}(z)^{2}.$$

Taking expectation over Y yields:

$$\mathbb{E}_{Z,Y}[\pi_n(Z)] = \mathbb{E}_Y \Big[\mathbb{E}_Z[\pi_n(Z) \mid Y] \Big] = \mathbb{E}_Y \left[\sum_{z \in S} \pi_n(z)^2 \right].$$

By Lemma C.6, we have:

$$\sum_{z \in \mathcal{S}} \pi(z \mid y)^2 \ge \exp(-H(\pi(\cdot \mid y))).$$

Averaging both sides over Y and applying Jensen's inequality (since $x \mapsto e^{-x}$ is convex), we obtain:

$$\mathbb{E}_{Z,Y}[\pi_n(Z)] = \mathbb{E}_Y \left[\sum_{z \in S} \pi(z \mid Y)^2 \right]$$

$$\geq \mathbb{E}_Y \left[\exp(-H(\pi(\cdot \mid Y))) \right]$$

$$\geq \exp(-\mathbb{E}_Y [H(\pi(\cdot \mid Y))])$$

$$= \exp(-H(Z \mid Y))$$

$$= \exp(-H_0 + I(Z; Y)),$$

where the last step uses the definition of mutual information.

Combining this with Equation (6) proves the claim.

Theorem C.5 shows that the success probability improves exponentially in the amount of mutual information gained through the lemmas, $I(Z; Y_{1:n})$. In particular, the success probability is strictly larger than in the case without lemmas, where $I(Z; Y_{1:n}) = 0$.

The following lemma was used in the proof of Theorem C.5:

Lemma C.6 (Relation Between Squared Sum and Entropy). For any probability distribution $p = (p_i)_i$, the following inequality holds:

$$\sum_{i} p_i^2 \ge \exp(-H(p)),$$

where $H(p) = -\sum_{i} p_i \log p_i$ denotes the Shannon entropy (with natural logarithm).

Proof. The log-sum inequality states that for nonnegative sequences $\{a_i\}, \{b_i\}$, the following holds:

$$\sum_{i} a_{i} \log \frac{a_{i}}{b_{i}} \geq \left(\sum_{i} a_{i}\right) \log \frac{\sum_{i} a_{i}}{\sum_{i} b_{i}}.$$

Let $a_i = p_i$ and $b_i = p_i^2$. Then the LHS becomes:

$$\sum_{i} p_{i} \log \frac{p_{i}}{p_{i}^{2}} = \sum_{i} p_{i} \log \frac{1}{p_{i}} = -\sum_{i} p_{i} \log p_{i} = H(p).$$

On the other hand, the right-hand side (RHS) becomes:

$$\left(\sum_{i} p_{i}\right) \log \frac{\sum_{i} p_{i}}{\sum_{i} p_{i}^{2}} = 1 \cdot \log \frac{1}{\sum_{i} p_{i}^{2}} = -\log \left(\sum_{i} p_{i}^{2}\right).$$

Hence, the log-sum inequality gives:

$$H(p) \ge -\log\left(\sum_{i} p_i^2\right).$$

Exponentiating both sides yields:

$$\sum_{i} p_i^2 \ge \exp(-H(p)).$$

This completes the proof.

D DETAILED EXPERIMENTAL SETUP

D.1 BENCHMARKING DATASET

We use the MiniF2F (Zheng et al., 2022) dataset, which consists of 488 mathematical problems formalized in Lean. These problems originate from sources such as AIME (American Invitational Mathematics Examination), AMC (American Mathematics Competitions), and IMO (International Math Olympiad) competitions, along with selected problems from the MATH dataset (Hendrycks et al., 2021), covering topics such as algebra, number theory, geometry, and analysis. Each problem is given as a Lean theorem statement. The benchmark is split into 244 validation and 244 test problems. We use the validation set during development (e.g., for tuning prompt formats) and report the final results on the test set. We use the revised version of miniF2F released by Wang et al. (2025); Ren et al. (2025).

Also, we observed that for problem names like algebra_2varlineareq_fp3zeq11_3tfm1m 5zeqn68_feqn10_zeq7, the LLM often struggled to reliably reproduce the latter part of the name due to its unintelligible character sequence. Therefore, we modified such problem names by removing the less interpretable suffixes and replacing them with simpler, more memorable labels such as algebra for our experiments.

D.2 USED MODELS

 For the informal LLM, we use <code>DeepSeek-R1-0528-Qwen3-8B</code> (DeepSeek-AI, 2025), a model obtained by distilling the chain-of-thought outputs of <code>DeepSeek-R1-0528</code> (DeepSeek-AI, 2025) into the Qwen3-8B (Yang et al., 2025a). This model surpasses Qwen3-8B on the AIME benchmark for natural language reasoning and achieves state-of-the-art performance at this scale. For the prover model, we use <code>Goedel-Prover-V2-7B6</code> (Lin et al., 2025c) and <code>DeepSeek-Prover-V2-7B7</code> (Ren et al., 2025), the state-of-the-art and second-best Lean 4 provers at this scale, respectively. For the formalizer model, we use <code>Goedel-Formalizer-V2-8B8</code> (Lin et al., 2025c) in the Goedel-Prover setup and <code>Kimina-Autoformalizer-7B9</code> (Wang et al., 2025). All of them are publicly available on Hugging Face (Wolf et al., 2020).

D.3 IMPLEMENTATION DETAILS

All models are invoked via vLLM (Kwon et al., 2023), a high-performance inference engine for large language models. We set max_num_batched_tokens and max_model_len parameters to 16384 to accommodate the long context lengths required for theorem proving, while keeping all other settings at their vLLM defaults. The models are run on NVIDIA A100 GPUs with 40GB of memory. We use Lean version 4.9.0 (Moura & Ullrich, 2021) throughout all experiments, following the same setup in Xin et al. (2025a); Ren et al. (2025); Lin et al. (2025c).

There are several bugs that may result in invalid Lean proofs being incorrectly accepted, such as the user-interference bug related to the <code>apply?</code> tactic discussed in version 2 of the arXiv paper by Ren et al. (2025), and a bug in REPL¹⁰. To avoid these issues and prevent invalid proofs from being mistakenly judged as correct, we check proofs with <code>lake build</code> instead of REPL and additionally verified that the <code>apply?</code> tactic is not used. Also, to avoid this bug and obtain reliable baseline results, we re-ran the experiments for Goedel-Prover-V2-8B. We used the official prompts provided on GitHub¹¹ and Hugging Face⁶, while keeping all other experimental settings strictly identical to those used in our method, thereby ensuring a fair comparison. For DeepSeek-Prover-V2, we relied on the results reported in version 2 of the arXiv paper (Ren et al., 2025), in which this bug has been fixed. All other baseline results are sourced from their respective papers.

D.4 SUMPLE BUDGET

We set $N_{\rm init}=N_{\rm refine}=50$. Thus, the sample budget at the initial direct proving stage is 50 at the first iteration, and 100 in total when including iterative refinement. For lemmas, we use $N_{\rm init}=N_{\rm refine}=10$ for each of the three lemmas. In the final synthesis stage, $N_{\rm init}=N_{\rm refine}=50$ is used again, resulting in a total sample budget of $50+50+(10+10)\times 3+50+50=260$.

D.5 BASELINE METHODS

We compare our approach against several baseline methods, categorized into two main classes: tree search methods and whole-proof generation methods. Tree search methods construct proofs incrementally by predicting individual tactics step by step, often guided by search algorithms such as best-first search or Monte Carlo Tree Search (MCTS). In contrast, whole-proof generation methods attempt to generate an entire proof script in a single forward pass, relying on the model's ability to plan the proof holistically.

The overview of the baseline methods used in our experiments is as follows:

Tree Search Method:

```
5https://huggingface.co/deepseek-ai/DeepSeek-R1-0528-Qwen3-8B
6https://huggingface.co/Goedel-LM/Goedel-Prover-V2-8B
7https://huggingface.co/deepseek-ai/DeepSeek-Prover-V2-7B
8https://huggingface.co/Goedel-LM/Goedel-Formalizer-V2-8B
9https://huggingface.co/AI-MO/Kimina-Autoformalizer-7B
10https://github.com/leanprover-community/repl/issues/44
11https://github.com/Goedel-LM/Goedel-Prover-V2
```

- DeepSeek-Prover-V1.5-RL + RMaxTS (Xin et al., 2025a) uses DeepSeek-Prover-V1.5-RL (Xin et al., 2025a), a 7B model trained with reinforcement learning, combined with RMaxTS (Xin et al., 2025a), a variant of MCTS that uses intrinsic rewards to explore diverse proof paths.
- InternLM2.5-StepProver + BFS + CG (Wu et al., 2024a) uses InternLM2.5-StepProver (Wu et al., 2024a), a 7B model trained via expert iteration (Anthony et al., 2017; Polu et al., 2023) starting with InternLM2-StepProver (Wu et al., 2024b), combined with a best-first search (BFS) strategy and a critic-guided (CG) sampling technique to explore longer proofs effectively.
- HunyuanProver v1.6 + BFS + DC (Li et al., 2025) uses HunyuanProver, a 7B model finetuned via a scalable data synthesis pipeline, in conjunction with best-first search guided by the distance critic (DC) to efficiently navigate complex Lean 4 proof search spaces.
- **BFS-Prover** (Xin et al., 2025b) uses a fine-tuned model of Qwen2.5-Math-7B model (Yang et al., 2024), trained through an expert-iteration pipeline. During inference, it employs a best-first search strategy to navigate the proof space efficiently.

Whole-Proof Generation Methods:

- Leanabell-Prover-GD-RL (Zhang et al., 2025) is a 7B model post-trained through continual training on statement-proof pairs and reinforcement learning using Lean 4 outcome rewards. This model is a fine-tuned version of Goedel-Prover-SFT (Lin et al., 2025b).
- **Goedel-Prover-SFT** (Lin et al., 2025b) is a 7B-parameter model obtained by supervised fine-tuning on DeepSeek-Prover-V1.5-Base (Xin et al., 2025a) with expert-iteration.
- STP: Self-Play Theorem Prover (Dong & Ma, 2025) employs a self-play framework that simultaneously takes on two roles, conjecturer and prover. The conjecturer is iteratively trained on statements that are barely provable by the current prover, incentivizing it to generate increasingly challenging conjectures. The prover uses standard expert iteration to verify and prove the generated conjectures. This model is a fine-tuned version of DeepSeek-Prover-V1.5-SFT (Xin et al., 2025a), which is a 7B-parameter model.
- **Kimina-Prover-Preview** (Wang et al., 2025) is a 72B-parameter reasoning model that learns specialized formal reasoning patterns via reinforcement learning. It is pretrained on a large corpus of formal proofs and fine-tuned with a binary correctness reward and consistency penalty. They also provide **Kimina-Prover-Preview-Distill-7B**, a distilled version from the 72B model.
- DeepSeek-Prover-V2 (Ren et al., 2025) uses DeepSeek-V3 to decompose each theorem into subgoals and then employs the proofs of those subgoals as cold-start data for reinforcement learning using binary correctness rewards and a consistency penalty to ensure that every subgoal appears in the final proof. It is implemented as a 671B-parameter model, and a distilled 7B-parameter variant is also provided.

E EXAMPLES OF SUCCESSFUL CASES ENABLED BY LEMMAS

E.1 DETAILED ANALYSIS

We analyze in detail the reasoning process for the problem induction_nfactltnexp nmlngt3, a case where the direct proof attempt failed but the use of auxiliary lemmas led to a successful proof. This problem asks for a formal proof that, for all natural numbers n>3, the inequality $n!< n^{n-1}$ always holds.

- The outputs for this problem, such as the generated lemmas, final formal proof, and the associated reasoning process, are provided in Appendix E.2 and after.
- In this case, the agent generated the following three lemmas: The first states that $3! < 3^{3-1}$; the second states that for any natural number $n \ge 2$, $n^{n-1} < (n+1)^{n-1}$; and the third states that for any natural number $n \ge 3$, $n! < (n+1)^{n-1}$. The first is a specific case of the original problem with n=3, while the second may provide a helpful hint toward solving the original problem. Both were easily proven in a single direct proof attempt. The third lemma generated in this case asserts that for

any natural number $n \ge 3$, $n! < (n+1)^{n-1}$. This lemma closely resembles the original problem, as it is a slightly weaker version of its conclusion. Due to its similarity and retained difficulty, the agent failed to construct a direct proof for it.

By examining the final successful reasoning trace in Appendix E.5, we see that the specific case for n=3, considered as the first lemma, appears explicitly on line 7. The reasoning also checks the cases for n=4 and n=5, following a similar pattern. Furthermore, as stated on line 13, the use of mathematical induction is clearly identified as the intended proof strategy. Then, the reasoning trace from line 14 to line 80 further elaborates the proof process within the framework of mathematical induction. Furthermore, in the final proof, the proof technique used in Lemma 2 is explicitly applied at lines 195–196.

Next, as a comparison, we analyze the reasoning process from the initial direct proving attempt without using any lemmas, as shown in Appendix E.6. Here, we present the reasoning trace that resulted in the fewest Lean errors among all initial direct attempts. Compared to the successful case with lemmas, we see that the proof strategy is much less clear in this direct attempt. In the "Key Observations" section (lines 6 to 14), there is no indication of using mathematical induction, unlike in the lemma-assisted case. Although the system explores several ideas from lines 15 to 63, the reasoning appears less focused and more exploratory, lacking a concrete plan. As a result, while it eventually leans toward using induction, the lack of a clear and structured approach prevents it from working out the necessary details, ultimately leading to failure in the formal proof, which tolerates no ambiguity.

This detailed case study highlights the effectiveness of our lemma-generation approach in uncovering viable proof strategies. This marks a significant advance over prior methods that decompose problems into subgoals, which often assume the overall proof strategy is known in advance. Identifying an initial proof strategy is often a challenging part of solving difficult problems. Indeed, Ren et al. (2025) employs a decomposition-based approach but relies on the much larger and stronger DeepSeek-V3 (DeepSeek-AI, 2024) to formulate the initial proof sketch. In contrast, our agent follows a reasoning process similar to that of human mathematicians when the proof strategy is not apparent at first glance, exploring specific cases or hypothesizing intermediate steps to discover a promising direction and ultimately uncover the overall proof strategy.

E.2 LEAN ENVIRONMENT SETUP

All Lean code was executed with the following header, following Xin et al. (2025a); Ren et al. (2025), which is omitted in the examples for brevity:

```
import Mathlib
import Aesop

set_option maxHeartbeats 0

open BigOperators Real Nat Topology Rat
```

Note that the following Lean code works correctly with Lean version 4.9.0, the environment used in our experiments, which strictly follows the same setup in Xin et al. (2025a); Ren et al. (2025). However, Nat.pow_lt_pow_of_lt_left has been deprecated in newer versions of Lean, so it must be replaced with Nat.pow_lt_pow_left to work with the latest environment.

E.3 THE PROBLEM

The following problem asks for a formal proof that, for all natural numbers n greater than 3, the inequality $n! < n^{n-1}$ holds:

E.4 THE GENERATED LEMMAS

The generated lemmas are shown below, each presented with both its original natural language form and the corresponding Lean formalization. The first and second lemmas were successfully proved and are presented with their proofs. The third lemma could not be proved and is therefore presented with a sorry placeholder.

```
1 ### Lemma 1: base_case_3
2 **Assumptions**:
3 None
4
5 **Conclusion**:
6 3! < 3^(3-1)</pre>
```

```
theorem base_case_3 : (3 : N)! < 3^(3 - 1) := by
have h_main : (3 : N)! < 3^(3 - 1) := by
-- Calculate the factorial and the power step-by-step
norm_num [Nat.factorial, Nat.pow_succ, Nat.mul_assoc]
-- Use 'decide' to confirm the inequality
<;> decide
exact h_main
```

```
1 ### Lemma 2: exponent_inequality
2 **Assumptions**:
3 n is a natural number and n ≥ 2
4
5 **Conclusion**:
6 n^(n-1) < (n+1)^(n-1)</pre>
```

```
theorem exponent_inequality \{n : \mathbb{N}\}\ (hn : 2 \le n):
      n^{(n-1)} < (n+1)^{(n-1)} := by
have h_main : n^{(n-1)} < (n+1)^{(n-1)} := by
 3
        have h_0 : n - 1 + 1 = n := by
         omega
 5
        have h_1 : n \geq 2 := by omega
 7
        have h_2: n-1 \ge 1 := by omega
        have h_3: n < n + 1 := by omega
 8
        -- Use the fact that if a < b and n > 0, then a^n < b^n
9
10
        exact calc
          n^{(n-1)} < (n+1)^{(n-1)} := by
11
             -- Apply the lemma that if a < b and n > 0, then a^n < b^n
12
             exact Nat.pow_lt_pow_of_lt_left h3 (by omega)
13
            = (n + 1)^{n} (n - 1) := by rfl
14
15
    exact h_main
```

```
1 ### Lemma 3: factorial_less_than_n_next_power
2 **Assumptions**:
3 n is a natural number and n ≥ 3
4
5 **Conclusion**:
6 n! < (n+1)^(n-1)</pre>
```

```
theorem factorial_lt_n_next_power \{n: \mathbb{N}\} (hn : 3 \leq n) : 
2    n ! < (n + 1) \hat{} (n - 1) := by sorry
```

E.5 THE FINAL PROOF AND ITS REASONING TRACE

We present below the final Lean proof, successfully constructed by taking the previously verified lemmas into account. Alongside the code, we also include the corresponding reasoning trace generated by the prover model that guided the overall proof construction.

```
1409
            ### Detailed Proof and Analysis
1410
         2
1411
           We are tasked with proving that for all natural numbers 'n \geq 3',
1412
               the factorial of 'n', i.e., 'n!', is less than 'n(n - 1)'.
1413
         4
           #### Key Observations:
         5
1414
           1. **Base Cases**:
         6
1415
               - For 'n = 3': '3! = 6' and '3^{\circ}(3 - 1) = 3^{\circ} = 9'. So, '6 < 9'
1416
               holds.
1417
               - For 'n = 4': '4! = 24' and '4^{\circ}(4 - 1) = 4^{3} = 64'. So, '24 <
1418
               64' holds.
1419
               - For 'n = 5': '5! = 120' and '5^{\circ}(5 - 1) = 5^4 = 625'. So, '120
1420
               < 625 holds.
               - The pattern suggests that this holds for all 'n \geq 3'.
1421
        11
1422
        12
           2. **General Proof Idea**:
1423
               - The proof can proceed by mathematical induction. We first
1424
               check the base case 'n = 3', and then for the inductive step,
1425
               we assume the statement holds for 'n = k' (i.e., 'k! < k^(k -
               1) ') and show that it also holds for 'n = k + 1' (i.e., '(k + 1')
1426
               1)! < (k + 1)^k.
1427
        14
               - The inductive step involves showing that (k + 1)! = (k + 1) *
1428
                k! < (k + 1) * k^{(k - 1)}. The last expression is (k + 1)^k
1429
               because k^(k-1) = k^k / k when multiplying by (k+1) is
1430
               not clean. Instead, we can directly use the fact that `(k +
               1)! = (k + 1) * k! and k! < k(k - 1) by the inductive
1431
               hypothesis, but we must also ensure k + 1 < k + 1 + 1 is not
1432
               relevant here. Wait, perhaps the best approach is to directly
1433
               manipulate the inequality (k + 1)! < (k + 1)^k.
1434
        15
1435
        16
               However, this seems messy. A simpler way is to compare '(k +
               1)!' and '(k + 1)^k' directly. For 'k \ge 3', we have '(k + 1)! =
1436
                 (k + 1) * k! < (k + 1) * k^k'.
1437
        17
1438
        18
               But this is not obviously true because 'k! < k^k' is equivalent
1439
               to k < k + 1, which is trivially true.
1440
        19
               Alternatively, (k + 1)! = (k + 1) * k! < (k + 1) * k^{(k - 1)}
        20
1441
               by IH, and (k + 1) * k^(k - 1) = k^k + k^(k - 1), which is
1442
               less than (k + 1)^k = k^k + k * k^(k - 1) + ..., but this is
1443
               getting complicated.
1444
        21
1445
        22
               A much better approach is to use induction directly. For 'n =
                3', we already checked this. Assume k! < k(k-1) for some
1446
                k \ge 3. Then:
1447
               (k + 1)! = (k + 1) * k! < (k + 1) * k^{(k - 1)}. Now, (k + 1) *
        23
1448
                k^{(k-1)} = k^k + k^{(k-1)}. We need this to be '< (k+1)^k
1449
                = k^k + k * k^(k - 1) + ...
1450
        24
               But 'k^k' cancels out, and we need 'k^k(k - 1) < k * k^k(k - 1) +
        25
1451
               ...', i.e., '1 < k + ...', which is true for 'k \geq 3'.
1452
        26
1453
        27
               However, the last step is not straightforward, and it might be
1454
                simpler to use induction concretely.
1455
        28
            #### Step 1: Abstract Plan
        29
1456
        30
1457
       31 1. **Base Case ('n = 3') **:
```

```
1458
               - Compute '3! = 6' and '3 ^{\circ} (3 - 1) = 9'. Since '6 < 9', the
1459
               base case holds.
        33
1460
            2. **Inductive Step ('n = k + 1')**:
        34
1461
               - Assume 'k! < k^{(k-1)}' holds for some 'k \geq 3'.
        35
1462
               - We need to show '(k + 1)! < (k + 1)^k'.
        36
               - '(k + 1)! = (k + 1) * k! < (k + 1) * k^(k - 1) ' by the
1463
        37
1464
               inductive hypothesis.
               - Next, we need to compare (k + 1) * k^(k - 1) to (k + 1)^k =
1465
                 (k + 1) * k^k / k.
1466
               - The inequality 'k * (k + 1) \le k * k + k * k = 2 * k^2' for 'k
        39
1467
                \geq 1' is too loose. A better approach is as follows:
1468
                  We want (k + 1) * k^{(k - 1)} < (k + 1)^k.
        40
                 - Divide both sides by (k + 1): k^{(k - 1)} < (k + 1)^{(k - 1)}
1469
        41
                1) `.
1470
                 - But:
        42
1471
        43
                   - This is false for k = 3: 3^2 = 9 < 4^2 = 16 holds.
1472
                For k = 4: 4^3 = 64 < 5^3 = 125 holds, etc. But we need
1473
                (k + 1) * k^{(k - 1)} < (k + 1)^{k}, \text{ not } k^{(k - 1)} < (k + 1)^{(k + 1)}
                - 1) `.
1474
               - Alternatively, perhaps a direct proof is better:
        44
1475
        45
                 - '(k + 1) * k^ (k - 1) = k^k + k^ (k - 1) '.
1476
                 - We must show k^k + k^(k - 1) < (k + 1)^k = k^k + k * k^(k
        46
1477
                - 1) + ... + 1 \.
1478
                 - Since k \ge 3, k^(k-1) < k * k^(k-1) and 0 < k^k
                (which cancels), we get k^(k-1) < k * k(k-1) because '1
1479
                < k' when 'k \geq 2', so 'k^k + k^(k - 1) < k^k + k * k^(k - 1) +
1480
                ...' is justified.
1481
        48
1482
            But the above is still messy. A better way is to prove (k + 1) *
        49
1483
                k^{k} - 1 < (k + 1)^k \text{ for } k \ge 3:
            (k + 1) * k^{k} - 1 = k^{k} + k^{k} - 1, and (k + 1)^{k} = k^{k} + k *
1484
                 k^{k} + 1 + ...
1485
        51 Subtract:
1486
            (k + 1)^k - (k + 1) * k^{k - 1} = k^k + k * k^{k - 1} + ... -
1487
                (k^k + k^{(k-1)}) = k * k^{(k-1)} + ... which is clearly
1488
                positive because all terms are positive.
           Alternatively, just note that k^k > 0 when k \geq 1 is sufficient
1489
        5.3
                because (k + 1)^k = k^k + ... \ge k^k > k^{k - 1}.
1490
        54
1491
            However, the direct approach is to compute (k + 1)^k = (k + 1) *
        55
1492
                ((k + 1) * ... * 1) \ge (k + 1) * (k * ... * 1) = (k + 1) * k^{(k - 1)}
1493
                1) ', which would work except for the last factor.
1494
        57
            Wait, we need (k + 1)^k > (k + 1) * k^{k - 1}. But (k + 1)^k =
1495
                (k + 1) * (k + 1)^{k} - 1.
1496
        5.8
1497
        59
            But (k + 1)^{k - 1} \ge k^{k - 1} is not immediately obvious (and
1498
                may even be false for 'k = 3' ('4^2 = 16 \geq 8 = 3^2' is true, but
1499
                ^{4}^{3} = 64 \ge 3^{3} = 27^{\prime} \text{ is true}).
1500
        60
           This seems correct, but we instead keep it simple:
1501
1502
        63 #### Inductive Proof Sketch:
1503
        1. Base Case ('n = 3'): '6 < 9' is true.
1504
        65 2. Inductive Step ('n = k + 1' assuming 'k! < k^(k - 1)'):
               - We have (k + 1)! < (k + 1) * k^{(k - 1)}.
1505
               - We need '(k + 1) * k^{(k - 1)} \le (k + 1)^k = (k + 1) * k^(k - 1)
        67
1506
               1) ' + ...
1507
               - Alternatively, the most straightforward induction is simply
1508
                to use that (k + 1) * k^(k - 1) < (k + 1)^k when k \ge 3.
1509
       70 This can be proved as follows:
1510
```

```
1512
        71 (k + 1)^k = (k + 1) * (k + 1)^{k - 1} and (k + 1)^{k - 1} \ge
                k^{k-1}' for 'k \geq 3' because:
1513
            (k + 1)^2 = k^2 + 2k + 1 \ge k^2 for all k \ge 1,
            (k + 1)^3 = k^3 + 3k^2 + 3k + 1 \ge k^3 + 3k^2 \ge k^3 + 2k^2 \ge k^3
1515
                 for k \geq 2, etc.
1516
        74 But we get (k + 1)^k \ge (k + 1) * k^{k - 1}, which is not what we
1517
                want! This is because we made a miscount. The correct relation
1518
                is:
            (k + 1)^k = (k + 1) * (k + 1)^{k - 1}., and (k + 1)^{k - 1} >
1519
                k^{k-1} for 'k \geq 3' (since 'l < m' implies 'l^{n} < m^{n}'
1520
                for 'l, m > 1'). That is, '(k + 1)^k > (k + 1) * k^{k-1} >
1521
                 (k + 1) * k^{k} - 1 - k^{k} - 1' is not true.
1522
            Instead, (k + 1)^k = (k + 1) * (k + 1)^{k - 1} \ge (k + 1) * k^{k - 1}
1523
                1}' because '(k + 1)^{k} - 1 \ge k^{k - 1}' for 'k \ge 3'.
1524
        78
1525
        79 Therefore, the full chain is:
1526
        80 (k + 1)! = (k + 1) * k! < (k + 1) * k^{(k - 1)} < (k + 1)^k.
1527
        81
        82 Thus, we can prove the inductive step using the above reasoning.
1528
        83
1529
        84
            #### Abstract Plan:
1530
        85
            1. Base Case:
1531
                - Verify the inequality for 'n = 3'.
        86
1532
        87 2. Inductive Step:
1533
                - Assume the inequality holds for n = k, i.e., k! < k(k - k)
        88
               1) ' for 'k \geq 3'.
1534
               - Prove for 'n = k + 1', i.e., '(k + 1)! < (k + 1)^k = (k + 1) *
        89
1535
                 k^{(k-1)}.
1536
                - But it is enough to prove (k + 1)^k \ge (k + 1) * k^(k - 1):
        90
                 - Since (k + 1)^k = (k + 1) * (k + 1)^{k - 1}, we must show
1537
                (k + 1)^{k} - 1 \ge k^{k} - 1.
1538
                 - This is true for 'k \geq 4' since '(k + 1)/k \geq k/(k - 1)'.
        92
1539
                 - Check 'k = 3' separately: '4^2 = 16 \ge 3^2 = 9' and '4^3 =
        93
1540
                 64 \ge 3^3 = 27, etc.
1541
        94
1542
        95 #### Lean 4 Proof Sketch with 'have' Statements
        96
1543
            '''lean4
        97
        98
            theorem induction
1545
              (n : \mathbb{N})
        99
1546
               (h_0 : 3 \le n) :
1547
        101
              (n) ! < n^{(n-1)} := by
               -- Base case: n = 3
1548
              have h_base : (3 : \mathbb{N})! < 3 ^ (3 - 1) := by sorry
        103
1549
              -- Inductive step: assume (k)! < k^{(k-1)} for some k \ge 3, prove
        104
1550
                (k + 1)! < (k + 1)^{(k + 1)} - 1)
1551
              have h_inductive_step : \forall (k : \mathbb{N}), 3 \leq k \rightarrow (k)! < k \hat{} (k - 1) \rightarrow (k + 1)! < (k + 1) \hat{} ((k + 1) - 1) := by sorry
        105
1552
              -- Combine base case and inductive step to get the general result
1553
        106
        107
             have h_{general}: (n)! < n ^ (n-1) := by sorry
1554
        108
              sorry
1555
            . . .
        109
1556
        110
1557
        111 This sets up the scaffolding of the proof using 'have' statements
1558
                with 'sorry'.
        112
1559
            ### Step-by-Step Abstract Plan
        113
1560
        114
1561
        115 1. **Base Case ('n = 3') **:
1562
               - Directly check '3! = 6' and '3^{\circ}(3 - 1) = 9', hence '6 < 9'.
       116
1563
       117
       118 2. **Inductive Hypothesis ('n = k') **:
1564
               - Assume 'k! < k^{(k-1)}' for some 'k \geq 3'.
       119
1565
```

```
1566
        3. **Inductive Step ('n = k + 1')**:
1567
                 - Prove (k + 1)! < (k + 1)^{(k + 1)} - 1) using:
                   - '(k + 1)! = (k + 1) * k! < (k + 1) * k^(k - 1) ' by IH.
1568
                  - '(k + 1) ^k = (k + 1) * (k + 1) ^{k - 1} \geq (k + 1) * k^(k - 1) ' if '(k + 1) ^{k - 1} \geq k^(k - 1) '. This holds for 'k \geq 3'
        124
1569
1570
                  because (k + 1) \ge k and functions are increasing.
1571
        125
1572
        126 4. **Conclusion**:
                 - By induction, the statement holds for all 'n \geq 3'.
        127
1573
        128
1574
        129 ### Lean 4 Proof Sketch with 'have' Statements and 'sorry's
1575
1576
             '''lean4
1577
        132
             theorem induction
                (n : \mathbb{N})
        133
1578
        134
                (h_0 : 3 \le n) :
1579
        135
                (n) ! < n^{(n-1)} := by
1580
               have h_base : (3 : \mathbb{N})! < 3 ^ (3 - 1) := by sorry
        136
1581
               have h_inductive_step : \forall (k : \mathbb{N}), 3 \leq k \rightarrow (k)! < k \hat{} (k - 1) \rightarrow
        137
                  (k + 1)! < (k + 1) ^ ((k + 1) - 1) := by sorry
1582
               have h_general : (n) ! < n ^ (n - 1) := by sorry
        138
1583
        139
               sorrv
1584
        140
1585
        141
1586
        142
             (Note: The 'have' statements are used to sketch the high-level
1587
                  structure of the proof. The actual Lean 4 proof would fill in
                  the 'sorry' placeholders with detailed reasoning.)
1588
        143
1589
        144 ### Complete Lean 4 Proof
1590
        145
1591
             '''lean4
        146
             theorem induction
1592
        147
                (n : \mathbb{N})
        148
1593
                (h_0 : 3 \le n) :
        149
1594
                (n) ! < n^{(n-1)} := by
        150
1595
               have h_base : (3 : \mathbb{N})! < 3 ^ (3 - 1) := by
        151
1596
        152
                norm_num
        153
                  <;> decide
1597
        154
               have h_inductive_step : \forall (k : \mathbb{N}), 3 \le k \to (k)! < k^{\hat{}} (k - 1) \to
        155
1599
                  (k + 1)! < (k + 1) ^ ((k + 1) - 1) := by
1600
                  intro k hk' hk''
        156
1601
        157
                  have h_1 : (k + 1)! = (k + 1) * k! := by
1602
        158
                   simp [Nat.factorial_succ, Nat.mul_comm]
        159
                    <;> ring
1603
        160
                    <;> simp_all [Nat.factorial]
1604
        161
                    <;> ring
1605
        162
                    <;> simp_all [Nat.factorial]
1606
        163
                  rw [h1]
                  have h_2: (k + 1 : \mathbb{N}) * k ! < (k + 1 : \mathbb{N}) * k ^ (k - 1) := by
1607
        164
                   -- Apply the inductive hypothesis to get the inequality for
        165
1608
                  k! < k \ (k - 1)
1609
                   have h_3 : k ! < k ^ (k - 1) := hk''
1610
        167
                    have h_4: (k + 1 : \mathbb{N}) > 0 := by linarith
1611
                    have h_5: (k + 1 : \mathbb{N}) * k ! < (k + 1 : \mathbb{N}) * k ^ (k - 1) := by
        168
                      exact Nat.mul_lt_mul_of_pos_left h3 (by linarith)
1612
        169
        170
                    exact h<sub>5</sub>
1613
                  have h_3: (k + 1 : \mathbb{N}) ^ ((k + 1) - 1) = (k + 1) ^ k := by
        171
1614
                    simp [Nat.sub\_add\_cancel (by linarith : 1 \le k + 1)]
        172
1615
        173
                    <;> ring_nf
1616
                    <;> simp_all
        174
        175
1617
                  rw [h_3]
        176
                  -- Use the fact that (k + 1) * k \hat{k} (k - 1) < (k + 1) \hat{k} to
1618
                  conclude the proof
1619
                 have h_4: (k + 1 : \mathbb{N}) * k \hat{(k - 1)} < (k + 1) \hat{k} := by
        177
```

```
1620
                178
                                       have h_5: k \ge 3 := by exact_mod_cast hk'
                                       have h_6: (k + 1 : \mathbb{N}) * k \hat{k} (k - 1) < (k + 1) \hat{k} := by
1621
                179
                                            have h_7: (k + 1 : \mathbb{N}) \hat{k} = (k + 1) * (k + 1) \hat{k} = (k + 1) * (k + 
                180
1622
1623
                181
                                                cases k with
1624
                                                 | zero => contradiction
                182
                                                 | succ k' =>
1625
                183
                                                     cases k' with
1626
                184
                                                     | zero => contradiction
                185
1627
                                                     | succ k'' =>
                186
1628
                                                         cases k'' with
                187
1629
                188
                                                          | zero => contradiction
                                                          | succ k''' =>
1630
                189
1631
                190
                                                              simp [Nat.pow_succ, Nat.mul_assoc]
                191
                                                              <;> ring_nf at *
1632
                                                               <;> omega
                192
1633
                193
                                           rw [h<sub>7</sub>]
1634
                                          have h_8: (k + 1 : \mathbb{N}) * k \hat{(k - 1)} < (k + 1) * ((k + 1) ^
                194
1635
                                   (k - 1)) := by
                                               have h_9: k ^ (k - 1) < (k + 1) ^ (k - 1) := by
1636
                195
                                                     exact Nat.pow_lt_pow_of_lt_left (by linarith) (by
                196
1637
                 197
                                                          have h_{10}: k - 1 \ge 2 := by
1638
                198
                                                            omega
1639
                199
                                                          omega)
1640
                200
                                                nlinarith
1641
                201
                                           nlinarith
                202
                                       exact h<sub>6</sub>
1642
                203
                                   nlinarith
1643
                204
1644
                              have h_general : (n)! < n ^ (n - 1) := by
                205
1645
                206
                                  have h_1: \forall n: \mathbb{N}, 3 \leq n \rightarrow (n)! < n \hat{ } (n-1) := by
1646
                207
                                       intro n hn
                                       induction' hn with k hk IH
                208
1647
                209
                                       \cdot -- Base case: k = 3
1648
                                          simp [Nat.factorial]
                210
1649
                211
                                           <;> norm_num
1650
                212
                                       · -- Inductive step: assume the statement holds for k, prove
                                  for k + 1
1651
                213
                                            cases k with
                214
                                            | zero => contradiction
1653
                                           | succ k' =>
                215
1654
                                                cases k' with
                216
1655
                217
                                                 | zero => contradiction
                                                | succ k'' =>
1656
                218
                                                    cases k'' with
                219
1657
                220
                                                     | zero => contradiction
1658
                                                     | succ k''' =>
                221
1659
                222
                                                          simp_all [Nat.factorial, Nat.pow_succ, Nat.mul_assoc]
1660
                223
                                                          -- Use the inductive hypothesis and the lemma
                                   'h_inductive_step' to prove the next step
1661
                224
                                                          <;>
1662
                225
                                                          (try omega) <;>
1663
                                                           (try nlinarith) <;>
                226
1664
                227
                                                           (try
1665
                228
                                                                  have h_2 := h_{inductive\_step} (k''' + 3) (by
1666
                229
                                   linarith) (by simpa [Nat.factorial, Nat.pow_succ,
1667
                                  Nat.mul_assoc] using IH)
1668
                                                                   simp_all [Nat.factorial, Nat.pow_succ,
                230
1669
                                  Nat.mul_assoc]
1670
                231
                232
                                                                   (try omega) <;>
1671
                233
                                                                    (try nlinarith)
1672
                234
                                                              })
1673
                235
                                                          <:>
```

```
1674
                                  (try
1675
         237
         238
                                       aesop
1676
         239
                                     })
1677
         240
                                  <;>
1678
         241
                                  (try
1679
         242
                                       nlinarith [pow_pos (by norm_num : (0 : \mathbb{N}) < 2)
1680
         243
                    k''', pow_pos (by norm_num : (0 : \mathbb{N}) < 2) (k''' + 1), pow_pos
1681
                    (by norm_num : (0 : \mathbb{N}) < 2) (k''' + 2)]
1682
         244
                                     })
1683
                    exact h<sub>1</sub> n h<sub>0</sub>
         245
1684
         246
1685
         247
                  exact h_general
               . . .
         248
1686
1687
```

E.6 COMPARISON: DIRECT PROOF ATTEMPT WITHOUT LEMMAS

1688

1689 1690

1691

1692

For comparison, we present below the reasoning process and generated Lean code from a direct proof attempt without using auxiliary lemmas. Among the 100 initial direct proof attempts ($N_{\rm init}=100$), we show the one that resulted in the fewest Lean errors.

```
1693
1694
            ### Detailed Proof and Analysis
1695
         3
           **Problem Analysis:**
1696
           We need to prove that for all integers 'n \geq 3', 'n! < n^(n - 1)'.
1697
1698
           **Key Observations:**
1699
           1. For 'n \geq 3', 'n - 1 \geq 2' (since '3 - 1 = 2'), so 'n^(n - 1)'
                grows faster than 'n!' as 'n' increases.
1700
            2. The factorial 'n!' is the product of all integers from '1' to
1701
                'n', while 'n^(n - 1)' is a product of 'n - 1' copies of 'n'
1702
                (assuming 'n > 3').
1703
            3. The inequality n! < n(n-1) can be compared term by term,
1704
               but a more precise estimate is needed.
1705
           4. For 'n \geq 3', we have 'n! < (n + 1 - 1)! = n!'? Wait, this is
        10
               not helpful. Instead, we can compare 'n!' with 'n^ (n - 1)':
1706
               - The inequality 'n! < n^{(n-1)}' is equivalent to '(n-1)! < n^{(n-1)}
        11
1707
               n^{(n-2)}, but this is not directly useful either.
1708
            5. A better idea is to directly compare 'n!' with 'n*(n
1709
               1) *...*(2) and n(n-1) = n * n * ... * n'.
               - For 'n \geq 3', each term 'k' in the product '2 * 3 * ... * (n -
1710
        1.3
                1) ' is '\geq 1' (except '2 * 1' is skipped if 'n \geq 4', but 'n \geq
1711
                3' is given), and each is '\leq n - 1' for 'k \leq n - 2'. However,
1712
                comparing 'k' to 'n' is complicated.
1713
               - Alternatively, 'n!' is the product of 'n' factors, each '\leq
        14
1714
                n', but 'n!' is a product of all numbers from '1' to 'n',
                whereas n^(n-1) is a product of n-1 copies of n.
1715
        15
1716
            But actually, induction is not straightforward here. The key is to
1717
                compare 'n!' with 'n * (n - 1) (n - 1)', since '(n - 1) (n - 1)
1718
                1) 'grows faster than '(n - 1)(n - 2) \dots' for 'n \ge 3'.
1719
        17
1720
           But 'n! < n * (n - 1)!' is 'n! < n * (n - 1)!', which is true
                except for 'n = 1'. So is 'n * (n - 1)! < n^n(n - 1)'? Not
1721
                obviously, because for 'n = 4', '4 \star 3! = 24' and '4<sup>3</sup> = 64',
1722
                but for 'n = 3', '3! = 6' and '3^2 = 9'.
1723
        19
1724
        20
           Instead, let's recall the definition of 'n!' and think about how
1725
               many times 'n!' is compared to 'n^(n-1)'.
        21
1726
           For 'n \geq 3', each time you compare 'n!' with 'n * (n - 1)!', and
1727
               this is '< n * (n - 1)! < n * n^(n - 2)  = n^(n - 1)  ? Seems
```

```
1728
               wrong because (n-1)! < n(n-2) would be true for n \ge 1
1729
               2' by induction, which is not a problem since we're proving it.
        2.3
1730
        24 Perhaps a better approach is to use telescoping inequalities.
1731
               Consider 'n! = n * (n - 1) * ... * 2 * 1', and 'n^ (n - 1)' is
1732
               the product of 'n - 1' copies of 'n'.
1733
        25
        26 For 'n > 3', each term in the factorial is smaller than 'n' ('k <
1734
               n - 1').
1735
1736
           But more precisely, n! = n * (n-1)! and (n-1)! < n^(n-2)
        28
1737
               for 'n \geq 3' (since 'n - 1 \geq 2').
1738
           Let's try to prove (n-1)! < n(n-2). If (n-1)! < n(n-1)!
1739
               2) \} holds, then 'n! = n * (n - 1)! < n * n^(n - 2) \} = n^(n -
1740
               2) } \.
1741
        31
1742
        32 But is 'n * n^(n - 2) = n^(n - 1)'? Yes, because 'n * n^(n -
1743
               2)} = n^{(1 + n - 2)} = n^{(n - 1)}.
1744
        33
           Thus, n! < n^(n-1), provided that (n-1)! < n^(n-2).
        34
1745
        35
1746
        36 Wait, 'n^(n - 2)}' is 'n * n^(n - 3)}' if 'n \geq 3'.
1747
        37
1748
        38 So the induction is:
        39 - Base case 'n = 3': '2! = 2' and '3^2 = 9', '2 < 9' holds.
1749
        40 - Inductive step: Assume for 'n', '(n - 1)! < n^n(n - 2)'. Then 'n
1750
               * (n-1)! < n * n^n(n-2) = n^n(n-1)'? But 'n * n^n(n-2)
1751
               = n^(n - 2 + 1) = n^(n - 1) '. So 'n! = n * (n - 1) ! < n^(n - 1) 
1752
               1)} \.
1753
        41
        42 This is a valid induction, provided that (n - 1)! < n^n(n - 2)
1754
               is true for 'n \geq 3'.
1755
        43
1756
        44 Now, recall that for 'n > 3':
1757
        45 'n! = n * (n - 1)!' and '(n - 1)! < n^n(n - 2)' if 'n \ge 3'.
1758
        46
        47 But is '(n - 1)! < n^(n - 2)}' true for 'n \ge 3'?
1759
        48
1760
           For 'n = 3', LHS = 2, RHS = 1. False. Hmm, '3! = 6' and '3'(3 -
        49
1761
               3) ' = 1? No, '3 \ge 3 and '3^(3 - 2) = 3^{1} = 3, '2 < 3,
1762
               so true.
1763
        50
        51 Wait no. For 'n = 4', '(3)! = 6', and '4\{4 - 2\}\} = 4^{2}\} = 16',
1764
               `6 < 16` true.
1765
        52
1766
           For 'n = 5', '(4)! = 24', '5^{6}5 - 2}} = 5^{3} = 125', '24 < 125'
        53
1767
               true.
1768
        55 Hypothesis that (n-1)! < n^n(n-2) for all n \ge 3.
1769
        56
1770
        57 Base case 'n = 3', '2! = 2 < 3^{1} = 3' true.
1771
1772
        Inductive step: Assume (n-1)! < n(n-2) for n \ge 4.
1773
        60
        61 We need 'n! < (n + 1)^n (n - 1)':
1774
            n! = n * (n - 1)! < n * n^(n - 2) = n^(n - 1) , since (n - 1)!
1775
               < n^(n - 2)} holds by the inductive hypothesis for n \ge 3.
1776
        63
1777
        64 This works!
1778
        65
           **Conclusion:** The inequality 'n! < n^{n-1}' holds for all 'n \geq
1779
               3' ('n >= 3').
1780
1781
```

```
1782
        68 But wait, the original Lean statement is 'n : \mathbb{N}' and '3 \leq n', so
1783
                indices are clear.
        69
1784
            We can proceed directly using induction. The 'have' statements
1785
                will reflect the proof.
1786
1787
            **Lean 4 Proof Sketch: **
        72
1788
        73
        74 1. **Base Case ('n = 3')**: Check that '3! < 3^{3} - 1' or '6 < 9'.
1789
        75 2. **Inductive Step**: Assume '(k)! < k^{k-1}' for some 'k \ge 3',
1790
                and prove (k + 1)! < (k + 1)^{k} using the assumption k! <
1791
                k^{k} \{ k - 1 \}'.
1792
1793
        77
            ### Abstract Step-by-Step Plan
        78
1794
        79 1. **Base Case ('n = 3') **:
1795
               - Compute 3! = 6.
1796
               - Compute 3^{3} - 1 = 3^{2} = 9.
1797
        82
               - Since '6 < 9', the statement holds.
1798
        8.3
           2. **Inductive Step ('n \rightarrow n + 1')**:
        84
                - Assume 'n! < n^{n-1}' for some 'n \geq 3' (induction
1800
               hypothesis).
1801
               - Prove (n + 1)! < (n + 1)^n:
1802
        87
                 - '(n + 1)! = (n + 1) * n!' (by definition).
                 - By hypothesis, n! < n^{n-1}.
1803
        88
                 - Thus, (n + 1)! < (n + 1) * n^{n - 1}.
        89
1804
        90
                 - It remains to prove (n + 1) * n^{n-1} \le (n + 1)^n,
1805
                i.e., n^{n-1} \le (n+1)^{n-1} (if n+1 \ge 2? Always
1806
                true since 'n \geq 2').
1807
                 * Since 'n \geq 2', 'n^{n - 1} \leq (n + 1)^{n - 1}' because 'n \leq
                n + 1' and the exponent is 'n - 1 \geq 2', making the term '(n +
1808
                1) ^{n - 1} strictly greater than ^n_{n - 1}.
1809
        92
1810
        93 ### Lean 4 'have' Statements
1811
        94
1812
           '''lean4
        95
        96 theorem induction
1813
              (n : \mathbb{N})
        97
        98
               (h_0 : 3 \le n) :
1815
              (n) ! < n^{n} (n - 1) := by
        99
1816
            have h_{main} : (n)! < n^{(n-1)} := by
1817
        101
                sorry
1818
              sorry
        103
1819
        104
1820
        105 ### Explanation of 'have' Statements:
1821
        106 1. The 'have' statement 'h_main' directly encodes our goal '(n)! <
1822
                n^{(n-1)}. Inside the proof, we would perform induction, but
                in a structured Lean proof, we break it into a single 'have
1823
                h main'.
1824
              - The actual proof would expand on 'h_main' by leveraging
1825
                induction, but for this sketch, it suffices to rely on the
1826
                good old induction hypothesis.
1827
        108
1828
        109 ### Complete Lean 4 Proof
        110
1829
            '''lean4
        111
1830
        112 theorem induction
1831
              (n : \mathbb{N})
        113
1832
        114
               (h_0 : 3 < n) :
               (n) ! < n^{(n-1)} := by
1833
       115
       116
              have h_{main} : (n)! < n^{(n-1)} := by
1834
                have h_1 : \forall n : \mathbb{N}, 3 \leq n \rightarrow (n)! \langle n \hat{n} - 1 \rangle := by
       117
1835
       118
               intro n hn
```

```
1836
                  induction' hn with n hn IH
        119
1837
                   · -- Base case: n = 3
                     norm_num [Nat.factorial]
1838
                   · -- Inductive step: assume the statement holds for n, prove
        122
1839
                 for n + 1
1840
        123
                      cases n with
1841
                      \mid zero \Rightarrow contradiction -- n cannot be zero since 3 < n
        124
                      | succ n =>
1842
        125
        126
                        cases n with
1843
                        | zero => contradiction -- n cannot be one since 3 \le n
        127
1844
        128
                        | succ n =>
1845
        129
                          cases n with
1846
                           | zero => contradiction -- n cannot be two since 3 \le n
        130
1847
        131
                           | succ n =>
                             simp_all [Nat.factorial, Nat.pow_succ, Nat.mul_comm,
        132
1848
                 Nat.mul_assoc, Nat.mul_left_comm]
1849
        133
                             <;>
1850
        134
                                (try omega) <;>
1851
        135
                                (try
1852
        136
                                   nlinarith [pow_pos (by decide : (0 : \mathbb{N}) < 2)
        137
1853
                 n, pow_pos (by decide : (0 : \mathbb{N}) < 2) (n - 1)]
1854
        138
                                 }) <;>
1855
        139
                                (try
1856
        140
1857
        141
                                    ring_nf at *
        142
1858
                                     nlinarith [pow_pos (by decide : (0 : \mathbb{N}) < 2)
        143
1859
                 n, pow_pos (by decide : (0 : \mathbb{N}) < 2) (n - 1)]
1860
        144
                                 }) <;>
1861
        145
                                (try
1862
        146
                                    simp_all [Nat.factorial_succ, Nat.mul_comm,
        147
1863
                 Nat.mul_assoc, Nat.mul_left_comm]
1864
        148
                                    <;>
1865
        149
                                      norm_num at *
1866
                                    <;>
                                      ring_nf at *
1867
        151
        152
1868
                                      nlinarith [pow_pos (by decide : (0 : \mathbb{N}) < 2)
        153
1869
                 n, pow_pos (by decide : (0 : \mathbb{N}) < 2) (n - 1)]
1870
        154
                                  })
1871
        155
                 exact h_1 n h_0
               exact h_main
1872
        156
        157
1873
1874
```