
BFS-Prover-V2: Scaling up Multi-Turn Off-Policy RL and Multi-Agent Tree Search for LLM Step-Provers

Anonymous Authors¹

Abstract

The integration of Large Language Models (LLMs) with automated theorem proving has shown immense promise, yet is constrained by challenges in scaling up both training-time reinforcement learning (RL) and inference-time compute. This paper introduces `BFS-Prover-V2`, a step-level theorem proving system designed to address this dual scaling problem. We present two primary innovations. The first is a novel multi-turn off-policy RL framework for continually improving the performance of the LLM step-prover at training time. This framework, inspired by the principles of AlphaZero, utilizes a multi-stage expert iteration pipeline featuring adaptive tactic-level data filtering and periodic retraining to surmount the performance plateaus that typically curtail long-term RL in LLM-based agents. The second innovation is a planner-enhanced multi-agent system that scales reasoning capabilities at inference time. This architecture employs a general reasoning model as a high-level planner to iteratively decompose complex theorems into a sequence of simpler subgoals. This hierarchical approach substantially reduces the search space, enabling a team of parallel prover agents to collaborate efficiently by leveraging a shared proof cache. We demonstrate that this dual approach to scaling yields state-of-the-art results on established formal mathematics benchmarks. `BFS-Prover-V2` achieves 95.08% and 41.4% on the miniF2F and ProofNet test sets respectively. While demonstrated in the domain of formal mathematics, the RL and inference techniques presented in this work are of broader interest and may be applied to other domains requiring long-horizon multi-turn reasoning and complex search.

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

1. Introduction

Automated Theorem Proving (ATP), a subfield of mathematical logic and automated reasoning, represents one of the foundational ambitions of computer science (Bibel et al., 1993). The contemporary landscape of formal mathematics is increasingly dominated by interactive theorem provers (ITPs) or proof assistants. These systems, such as Coq, Isabelle, and Lean, require a human user to guide the proof process, but they automate significant deductive tasks and, most importantly, provide a machine-checkable guarantee of correctness (Geuvers, 2009). Among these, the Lean 4 programming language (Moura & Ullrich, 2021) has emerged as a particularly vibrant ecosystem. A key factor in its success is Mathlib (Blokpoel, 2024), a vast, comprehensive, and community-driven library of formalized mathematics. Spanning over a million lines of code, Mathlib covers extensive areas of algebra, analysis, topology, and more, providing a rich foundation for both advanced mathematical research and the development of verified systems.

The rise of Lean 4 has coincided with the explosion in the capabilities of LLMs (OpenAI, 2023; Comanici et al., 2025; Seed et al., 2025), opening a new frontier in neuro-symbolic AI systems. The goal here is to integrate the intuitive yet powerful generation and search capabilities of LLMs with the absolute logical verification of formal systems. This research direction centers on a key feedback loop: an LLM proposes intuitive proof steps, the Lean compiler provides rigorous verification, and RL (Sutton, 2018) uses that verification to continuously improve the LLM’s reasoning abilities (Yang et al., 2024b; Xin et al., 2024a; Polu et al., 2022; Han et al., 2021; Lample et al., 2022).

1.1. Motivation: A Duality of Scaling Challenges in Reasoning

The development of high-performance formal math provers, or any other reasoning agents, is contingent upon solving two fundamental and deeply interconnected scaling challenges.

Training-time scaling. This refers to the techniques required to continuously enhance a model’s foundational capabilities and tactical intuitions via training. A common

and significant obstacle in applying RL to LLMs is the phenomenon of performance plateaus: after an initial phase of rapid improvement, models often stagnate, with their capabilities ceasing to grow despite continued training (Liu et al., 2025; Team et al., 2025; Yu et al., 2025; Yue et al., 2025; Guo et al., 2025; Seed et al., 2025; Xin et al., 2024a;b). Overcoming this limitation requires carefully designed algorithms that can sustain learning over extended periods, enabling the model to transition from mastering simple problems to tackling increasingly complex theorems.

Inference-time scaling. This addresses the method of using a trained model to solve previously unseen theorems. The primary bottleneck here is inefficient exploration in a vast search space. Specifically, pure tree search is often constrained by a lack of global planning ability and a search space that grows exponentially with proof depth. More generally, current inference paradigms suffer from the fact that search trajectories are independent, meaning that insights gained in one proof attempt are not shared with others. The challenge, therefore, is to design an inference architecture that incorporates planning and collaborative search to more effectively allocate computational resources (Baba et al., 2025; Zhou et al., 2025; Chen et al., 2025b; Jiang et al., 2022; Cao et al., 2025).

1.2. Our Contributions

This paper presents BFS-Prover-V2, a comprehensive training and inference system for neural theorem proving in Lean 4 that introduces novel solutions to the above scaling challenges. The primary contributions of this work are as follows:

Novel RL Scaling Techniques at Training: We develop a distillation-free multi-stage expert-iteration framework (Silver et al., 2018; Anthony et al., 2017), a form of off-policy RL, tailored for the domain of formal theorem proving. To sustain learning and overcome performance plateaus, we introduce a suite of specialized techniques within the RL pipeline. These include an adaptive, perplexity-based data filtering strategy at the tactic level, which creates an automated curriculum for the agent, and a periodic retraining mechanism that acts as a “soft reset” to escape local optima in the model parameter space and increase model scaling potential.

Multi-Agent Tree Search System at Inference: For inference-time scaling, we introduce a hierarchical reasoning architecture. A general-purpose reasoning model, termed the planner, provides global planning ability by iteratively decomposing complex theorems (or goals) into a sequence of more manageable subgoals. These subgoals serve as tree search “checkpoints” with successful proof trajectories stored in a shared cache. This dramatically reduces search complexity by converting the total computational

effort from a product of individual subgoal complexities to their sum.

State-of-the-Art Empirical Results: We validate the effectiveness and generalizability of our dual scaling approach on established benchmarks. In particular, BFS-Prover-V2 achieves 95.08% on the miniF2F test set, largely surpassing previous LLM step-provers (Wu et al., 2024a; Xin et al., 2025; Liang et al., 2025) and performing on par with the best whole proof generation models (Ren et al., 2025; Lin et al., 2025b; Wang et al., 2025). On ProofNet-test, it achieves 41.4%, setting a new state-of-the-art and demonstrating robust generalization across distributions.

2. The BFS-Prover-V2 System

This section details the two core components of BFS-Prover-V2: (i) a training pipeline, grounded in a Markov Decision Process (MDP) (Puterman, 1990) and scaled via adaptive filtering and periodic retraining; and (ii) an inference engine, which uses a planner-enhanced multi-agent search for hierarchical reasoning. These components build upon the foundation of BFS-Prover-V1 (Xin et al., 2025) to specifically address the dual challenges of scaling at both training and inference time. We provide visual overviews of these components in Fig. 1 and 3, with practical implementation parameters and ablations detailed in Section 3.

2.1. A Step-Level Formulation: Theorem Proving as a Markov Decision Process

We formulate proof search in Lean 4 tactic mode as a Markov Decision Process (MDP), where an LLM-based prover (agent) interacts with the Lean proof checker (environment) to construct formal proofs. This formulation naturally captures the sequential, stateful nature of tactic-based formal theorem proving. Formally, we define the MDP tuple $\mathcal{M} = (S, A, P, R)$ as follows:

- **State Space S :** Each state $s \in S$ corresponds to a tactic state returned by the Lean compiler, comprising the current hypotheses (known facts) and target goals to be proven.
- **Action Space A :** An action $a \in A$ is a tactic string generated by the prover. Each tactic encodes a proof step, such as theorem application, term rewriting, or goal decomposition, that instructs the compiler to perform a specific deductive transformation.
- **Transition P :** The transition function $P(s' | s, a)$ is deterministically executed by the Lean proof checker. Given state s and action a , the compiler either produces a successor state s' if a is applicable, or returns an error, resulting in a terminal failure state.

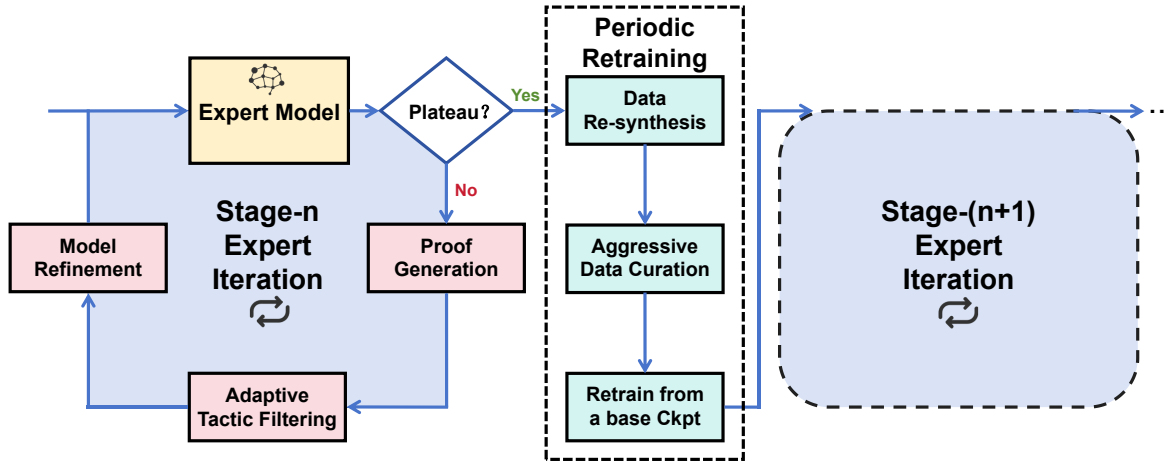


Figure 1. Overview of the training-time scaling up architecture. The process begins with a current expert model. The system then evaluates the model’s performance to check for a plateau. If performance is improving, the model enters an inner **expert iteration loop**. Conversely, if performance has plateaued, the system triggers the outer **retraining loop**. The retraining loop yields an improved expert model serving as the starting point for the next cycle.

- **Reward Function R :** We employ sparse rewards where $R(s, a) = 1$ if and only if the state-action pair (s, a) lies on a trajectory culminating in a successful proof. Otherwise, $R(s, a) = 0$.

This tactic-level stepwise interactive formulation fundamentally differs from whole-proof generation approaches (Ren et al., 2025; Lin et al., 2025b; Wang et al., 2025), which frame theorem proving as a one-shot code generation task from problem statements to complete proof scripts. While simpler, such approaches cannot adapt to intermediate proof states and lack integration with interactive theorem proving workflows (Welleck & Saha, 2023; Song et al., 2024). Our MDP-based approach, by design, trains an agent that functions as a genuine Lean copilot, capable of suggesting appropriate tactics at each step in the proof process (Yang et al., 2024c).

2.2. Scaling up Training: Multi-Stage Expert Iteration

The core training loop of BFS-Prover-V2 is an expert iteration pipeline, which may be viewed as a variant of the reinforcement learning algorithm used in AlphaZero (Anthony et al., 2017; Silver et al., 2018). This approach enables the system to learn and improve its theorem-proving capabilities from its own experience (Silver & Sutton, 2025). The process, illustrated in the inner loop of Fig. 1, includes two major alternating phases: proof generation and model refinement.

Phase 1: Proof Generation: The current best version of the LLM prover, referred to as the expert, is tasked with solving a large corpus of mathematical problems. The expert model is coupled with the best-first tree search (BFS) algorithm used in BFS-Prover-V1 (Xin et al., 2025) to explore the

vast space of possible proof trajectories. Each successful proof found during this phase constitutes a trajectory of (state, tactic) pairs. Across a single round of expert iteration, the system performs approximately 10^7 tree searches, generating a massive synthetic dataset.

Phase 2: Model Refinement: The state-tactic pairs from the successful proof trajectories generated in the first phase are used to update the model’s parameters. The updated model then becomes the new “expert” for the next round of iteration.

A central challenge in scaling the expert iteration pipeline or RL in general is managing the vast quantity and variable quality of the generated trajectories. Naively training on every successful trajectory quickly leads to diminishing returns, performance stagnation, and mode collapse (Liu et al., 2025; Sutton, 2018; Xin et al., 2025). To sustain improvement over many iterations, we introduce two key algorithmic innovations: a dynamic, fine-grained data filtering strategy and a periodic full-model retraining process. These techniques work in concert to form an automated curriculum that continuously improves the agent’s capability over a long horizon. The overall architecture of this pipeline is illustrated in Fig. 1, and we detail each of these innovations in the following subsections.

2.2.1. ADAPTIVE TACTIC FILTERING: LEARNING FROM THE “JUST RIGHT” DATA

Instead of relying on coarse, problem-level filtering (Yu et al., 2025; Team et al., 2025), which often uses static metrics of difficulty, we adopt a more fine-grained approach at the tactic level. This strategy is guided by the empirical observation that the perplexity (negative log-probability)

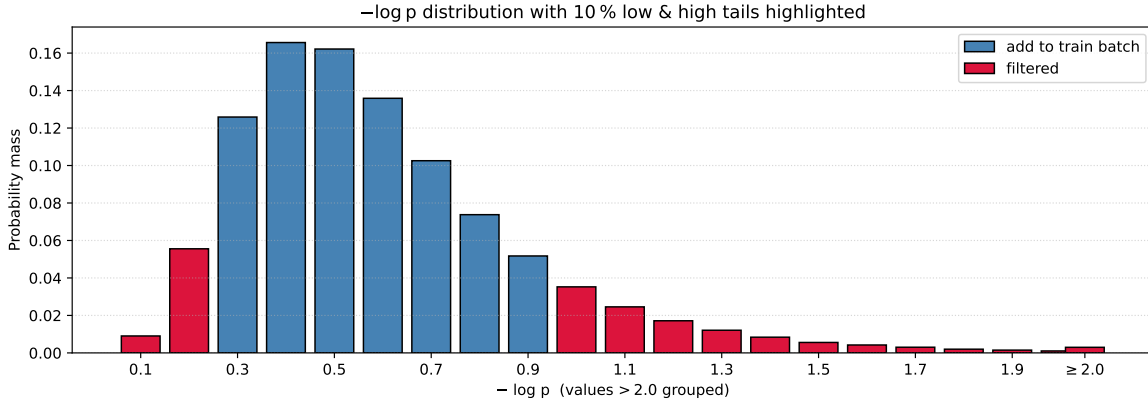


Figure 2. Tactic-Level Data Filtering Based on the Perplexity Distribution. This histogram shows the probability distribution of tactic perplexity (represented as normalized negative log-probability) from a single round of expert iteration. We filter out the low- and high-perplexity tails, shown in red.

of tactics generated by the LLM roughly follows a Gaussian distribution. The distribution, shown in Fig. 2, can be divided into three distinct regions, each with different implications for learning:

- **The Low-Perplexity Tail:** This region corresponds to tactics for which the model has very high confidence. These are typically simple, “obvious” steps, such as basic simplification or applying a clear-cut hypothesis. Including these examples in the training batch offers no new learning signal; it merely reinforces what the model already knows well and can contribute to overfitting and a reduction in exploratory capacity.
- **The High-Perplexity Tail:** This region represents tactics that the model finds highly surprising. Our case studies reveal that these are often not instances of brilliant, non-obvious reasoning. Instead, they frequently correspond to noisy or suboptimal choices, such as using a powerful, general-purpose tactic with many unnecessary parameters on a simple problem where a more direct tactic would suffice. These “fancy” operations can be detrimental to training, as they may teach the model to generate overly complex or irrelevant tactics, leading to hallucinations and degrading its core reasoning ability.
- **The Central Distribution:** This is the “goldilocks” zone. The tactics in this region are neither too easy nor too noisy. They represent steps that are challenging for the model but still within its grasp. By selectively training only on the data from this central part of the distribution, we ensure that the model is constantly learning at the edge of its capabilities.

This adaptive filtering mechanism functions as a fully automated form of curriculum learning. It does not rely on any

external (or predefined) metric of difficulty. Instead, it uses the model’s own uncertainty (as measured by perplexity) as a dynamic signal of what constitutes valuable training data at its current stage of development. This ensures a smooth and stable evolution of the model’s internal policy distribution throughout training, enabling sustained growth in performance. To further illustrate this filtering strategy, we present real-world examples of tactics falling into these three categories.

1. Low-Perplexity Tail

Trivial tactics (**Discard**)

```
x: ℝ
h0: x ∈ Set.Icc (π/2) (3π/2)
h1: 0 ≤ cos x
h2: ⊢ π/2 ≤ x ∧ x ≤ 3π/2
```

exact h0

2. High-Perplexity Tail

Noisy/hallucinated tactics (**Discard**)

```
p: ℝ
h0: 0 < p
hp: ¬p = 0
h2: (800+5*p) * (7*p) = 800*10*p
h3: ⊢ 7*p = 48*10
```

```
linarith [mul_pos h0 (show 0 < p by linarith),
mul_pos h0 (by linarith: 0 < p/2), mul_pos h0 (
by linarith : 0 < p)]
```

3. Mid-Perplexity Zone

Informative tactics (**Keep**)

```
x: ℝ
hx: x = 250000
h3: ⊢ x = 2.5 * 10^5
```

norm_num [hx]

2.2.2. PERIODIC RETRAINING: A “SOFT RESET” TO ESCAPE LOCAL OPTIMA

Even with adaptive filtering, performance eventually plateaus as the prover’s tactic preferences become increasingly rigid, causing it to overfit to a narrow set of proof

patterns and settle into a local optimum. It becomes very good at solving problems in particular ways, but loses the ability to discover novel approaches required for harder or new problems. To escape local optima and reinvigorate the learning process, we introduce a periodic “soft reset” procedure. This constitutes a multi-stage expert-iteration process designed to increase the model’s entropy and reset its exploratory potential without losing the competence it has already gained. The procedure is as follows:

1. **Re-synthesis and De-noise:** We re-run the current expert prover on the full historical problem set to regenerate all proofs using its improved policy. Since the prover at this stage is substantially stronger than the checkpoints that produced the earlier trajectories, the newly synthesized proofs are typically shorter, structurally cleaner, and contain fewer spurious steps. This pass serves as a model-driven denoising phase: it replaces outdated trajectories with higher-quality ones and removes redundant or circuitous reasoning patterns that accumulated during earlier, less capable iterations.
2. **Aggressive Data Curation:** The new, higher-quality proofs generated in the data re-synthesis phase are then subjected to an aggressive version of the tactic-level perplexity filtering described above. A much larger portion of the data is discarded, retaining only the most crucial and informative tactic steps.
3. **Retrain from a base Checkpoint:** The existing training data is completely replaced by this new, highly curated, and compact dataset. A fresh model instance is then initialized from the base checkpoint and trained from scratch on this refined data.

The resulting model initially exhibits a temporary drop in performance. This is expected, as it has been trained on a smaller, more focused dataset and has forgotten some of its previous stylistic biases. However, this new model possesses a significantly higher exploratory potential. When it is reintroduced into the expert iteration loop, its increased capacity for exploration allows it to discover novel proof strategies that were inaccessible to the previous, over-specialized model. Consequently, its performance rapidly recovers and then surpasses the previous peak.

2.3. Scaling up Inference: Planner-Enhanced Multi-Agent Search

Many complex proofs in mathematics are not found through a linear sequence of simple deductions but rather through a hierarchical process of identifying and proving crucial intermediate results. Likewise, we introduce a hierarchical inference architecture, shown in Fig. 3, that divides the

labor of theorem proving between two distinct agents: a high-level planner and a low-level prover.

Planner: This is a general-purpose reasoning LLM tasked with goal decomposition. Given the current theorem statement and proof progress, its role is not to generate a specific tactic but to propose a high-level plan that includes a series of intermediate subgoals. By formulating these subgoals, the planner effectively transforms a single, monolithic, and potentially intractable search problem into a structured sequence of smaller, more manageable ones. This decomposition substantially reduces the dimensionality of the search space that the prover must explore.

Prover: This is the specialized LLM tactic generator trained via our multi-stage expert iteration pipeline described in Section 2.2. It receives one subgoal at a time from the planner and uses its learned policy, in conjunction with Best-first tree search algorithm (Xin et al., 2025), to find a formal proof for that specific subgoal.

This division of labor mirrors the cognitive workflow of a human mathematician, who might first sketch out the high-level structure of a proof by identifying key lemmas (the planner’s role) and then proceed to work out the detailed, step-by-step deductions for each lemma (the prover’s role). This hierarchical structure is a powerful architectural pattern for tackling complex reasoning tasks.

2.3.1. OPERATIONAL MECHANICS OF PLANNER-GUIDED SEARCH

As shown in Fig. 3, the interaction between the planner and the prover system unfolds in a dynamic loop, allowing the plan to be revised as proof search progresses:

1. **Initial Planning:** At the start of a proof attempt, the planner is queried with the main theorem statement. It returns a list of proposed subgoals, formatted as Lean `have` statements.
2. **Sequential Proof of Subgoals:** The prover system addresses the subgoals one by one. It takes the first subgoal in the queue and initiates tree search to find its proof.
3. **Context Augmentation:** When a subgoal is solved, its statement is incorporated into the proof context and becomes available as a hypothesis for all remaining subgoals and the main theorem itself.
4. **Dynamic Replanning:** If the prover exhausts its search budget on a subgoal, which may occur either because the subgoal is intrinsically difficult or because the planner previously produced an incorrect subgoal, the system does not terminate. Instead, the planner is invoked again with an augmented input containing the

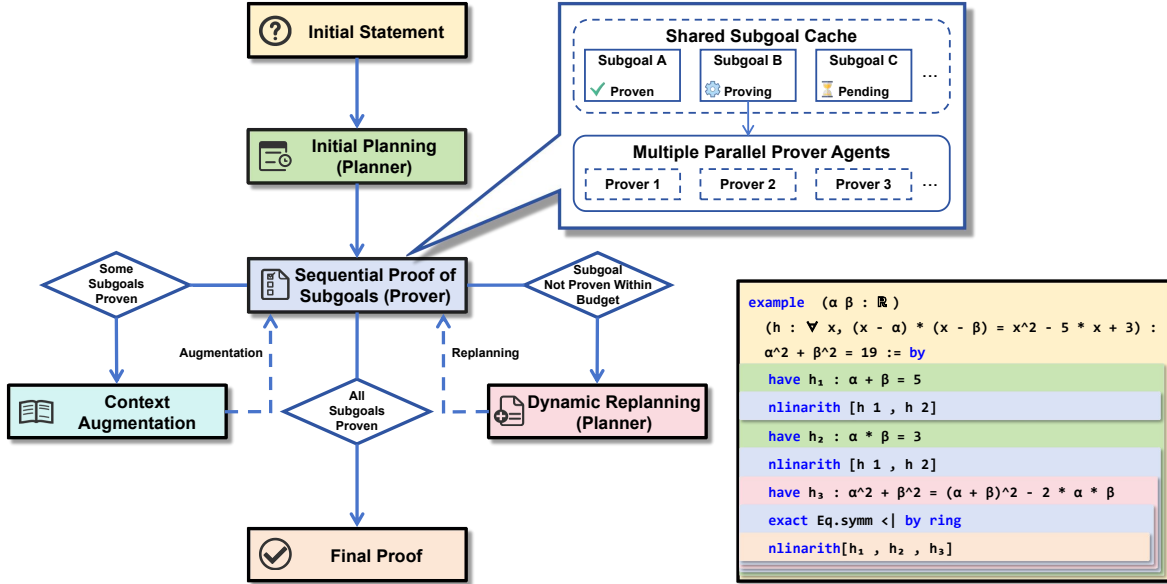


Figure 3. Overview of the multi-agent tree search architecture. The **Planner** agent decomposes the main theorem into a sequence of simpler subgoals, which are managed in a **Shared Subgoal Cache** and solved in parallel by multiple **Prover** agents. Successfully proven subgoals augment the main proof’s context, while failures trigger a **Dynamic Replanning** loop. The inset provides an example, demonstrating how proving intermediate lemmas (h_1, h_2, h_3) facilitates the proof of the final goal.

theorem statement and all previously proven subgoals. The planner then generates a revised decomposition that typically corrects, refines, or further subdivides the subgoal where search stalled.

This dynamic and iterative loop between planning and proving makes the **BFS-Prover-V2** system resilient to getting stuck, effectively scaling its inference-time capabilities to tackle complex theorems that would be intractable for a monolithic tree search.

2.3.2. MULTI-AGENT COLLABORATION VIA FOCUSED PARALLELISM AND SHARED SUBGOAL CACHE

To further scale inference-time compute, we deploy the planner–prover architecture in a multi-agent setting in which several prover instances run in parallel. These agents jointly execute the subgoal sequence proposed by the planner, coordinated through two design principles: focused parallelism and a shared subgoal cache.

Focused parallelism: Rather than distributing different subgoals in parallel across agents, all prover instances allocate their full search budget to a single active subgoal before the system advances to the next. This sequential execution concentrates search effort on difficult reasoning bottlenecks where only a subset of provers would need substantially more time to progress and avoids wasted compute on downstream subgoals that would be invalidated if an earlier step fails and triggers a replan.

Shared Subgoal Cache: This cache is the central commu-

nication and state-tracking mechanism, shared across all parallel prover instances. It stores the full sequence of subgoals generated by the planner, tracks the real-time status of each subgoal (e.g., pending, proving, proven), and records the proof for any solved subgoal.

This architecture creates a cooperative sprint for each lemma in the plan. When a new subgoal is designated as the active target, all prover agents begin independent tree searches for that single subgoal in parallel. As soon as the first agent finds a valid proof, it writes the result to the shared cache. The subgoal cache signals all other agents to terminate their search, preventing redundant computation. The entire group of agents then proceeds to the next subgoal in the sequence.

3. Experiments

3.1. Practical Implementation

Base model and Data: Our prover is built upon Qwen2.5-Math-7B and Qwen2.5-32B (Yang et al., 2024a), which serve as the base for our policy optimization. The multi-stage expert iteration process was initialized from the checkpoint **BFS-Prover-V1** (Xin et al., 2025). To construct a large-scale training corpus, we autoformalized the NuminaMath-CoT and NuminaMath-1.5 datasets (Li et al., 2024) by prompting general-purpose LLMs, augmented with Lean 4 compiler feedback. Together with Goedel-Pset-V1 (Lin et al., 2025a), this yields roughly 3 million formal statements (Wu et al., 2024b; Ying et al., 2024; Blokpoel, 2024). Prompts used for autoformalization can be found in

Section C.1. All experiments are conducted in Lean v4.10.0 with LeanDojo (Yang et al., 2024c).

Training setup: After each expert iteration round, we refine the policy LLM using one of two supervised fine-tuning (SFT) strategies, selected based on the outcome of the round. In the inner expert-iteration loop, we fine-tune the current best checkpoint for one epoch with cosine learning rate decay from 5×10^{-6} to 1×10^{-7} . When periodic retraining is triggered (Section 2.2), we train for three epochs with a larger learning rate that decays from 2×10^{-5} to 1×10^{-6} . Both strategies use a global batch size of 1024.

Inference configuration: Our inference pipeline combines a low-level prover with a high-level planner, as detailed in Section 2.3. Prover agents perform best-first search (BFS) following the BFS-Prover-V1 implementation (Xin et al., 2025). Unless stated otherwise, we use a sampling temperature of 1.3, an expansion width of 3, and a length-normalization factor of 2.0 during expert iterations. For the planner, we use Gemini 2.5 Pro; other general-purpose reasoning models can reach similar performance given suitable prompts. Planner prompts are provided in Section C.2.

3.2. Benchmark Results

We evaluated BFS-Prover-V2 on two primary benchmarks: miniF2F (Zheng et al., 2021), which targets high-school mathematical Olympiad problems (in-distribution), and ProofNet (Azerbayev et al., 2023), which covers undergraduate textbook level mathematics and serves as a rigorous test for out-of-distribution (OOD) generalization.

As detailed in Table 1, our system establishes a new state of the art among step-level tree-search provers. On the miniF2F-test set, our 32B model with the planner achieves an accuracy of 95.08% (95.49% on miniF2F-valid), while the 7B model reaches 92.6%.

Crucially, our approach demonstrates robust OOD generalization on ProofNet-test. The 32B model achieves 41.4%, and the 7B model reaches 34.4% via pure tree search, notably surpassing the 671B DeepSeek-Prover-V2 (37.1%) despite being 20 times smaller, and its 7B variant (29.6%), respectively. We attribute this superior OOD performance to the inherent flexibility of step-level proving: unlike whole-proof generation models that rely heavily on the training distribution, a trained step-prover can adapt its exploration strategy by adjusting search parameters at test time to match problem distributions, enabling effective transfer without retraining.

3.3. Further Analysis on Training

We report training-time ablations that justify the utility of tactic-level data curation, the critical role of periodic retrain-

Table 1. Comparison with other leading theorem provers. † denotes concurrent work.

Prover Method	Budget	miniF2F ProofNet	
<i>Tree-search provers</i>			
InternLM2.5-Step-7B	$256 \times 32 \times 600$	65.9%	$\approx 27\%$
Hunyuan-Prover-7B	$600 \times 8 \times 400$	68.4%	-
BFS-Prover-V1-7B	$2048 \times 2 \times 600$	70.8%	-
	accumulative	73.0%	-
MPS-Prover-7B†	$64 \times 4 \times 800 \times 8$	72.5%	-
	accumulative	75.8%	-
BFS-Prover-V2-7B	accumulative	82.4%	34.4%
w/ Planner	accumulative	92.6%	-
BFS-Prover-V2-32B	accumulative	86.1%	41.4%
w/ Planner	accumulative	95.1%	-
<i>Whole-proof provers</i>			
DeepSeek-Prover-V2-7B	8192 / 1024	82.0%	29.6%
w/ Prover Agent	260	82.8%	-
DeepSeek-Prover-V2-671B	8192 / 1024	88.9%	37.1%
Kimina-Prover-72B†	1024	87.7%	-
w/ TTRL search	accumulative	92.2%	-
Goedel-Prover-V2-7B†	8192	90.2%	-
w/ Prover Agent	260	86.5%	-
Goedel-Prover-V2-32B†	8192	92.2%	-
w/ Self-correction	1024	92.6%	-
Delta-Prover†	accumulative	95.9%	-
Seed-Prover†	accumulative	99.6%	-

ing in escaping local optima, and the motivation for scaling to larger base models.

Perplexity-based tactic filtering: We investigated the impact of our filtering strategy during the multi-stage expert iterations. We conducted an ablation on Checkpoint 3 in Fig. 4 (derived after the first retraining phase). The training corpus consisted of 459,540 pairs of human data. Without tactic filtering, the combination of human data and synthetic expert iteration data yielded 857,897 state-action pairs. Training on this unfiltered set resulted in a validation loss of 0.5597 and a miniF2F test score of 75%. In contrast, applying our perplexity-based filtering reduced the dataset to 660,254 high-quality samples. Despite the smaller data volume, this filtered run (Checkpoint 4) resulted in a higher validation loss (0.6211) but a superior test score of 76.63%. Further validation on checkpoint 4 confirmed the approach’s effectiveness: training with all expert iteration data resulted in performance degradation to 75.81%, while continued filtering improved performance to 77.04%.

Periodic retraining: Performance plateaus appeared at several checkpoints during training, where continued expert iteration did not improve and sometimes even reduced performance. Our periodic retraining mechanism consistently overcame these local optima. At checkpoint 2, performance stagnated at 75.41 percent. After retraining, accuracy briefly decreased to 75.00% but then increased to 76.64% in the




Figure 4. Sustained Performance Improvement through Expert Iteration and Periodic Retraining. This graph plots the prover’s performance on the miniF2F benchmark vs. the expert iteration rounds.

next iteration. A similar pattern occurred at checkpoint 6: accuracy was 78.28% before retraining, declined to 77.05% immediately after retraining, and then reached 79.92% after two additional iterations. Crucially, this phenomenon is scalable: the larger 32B model exhibited a similar plateau at Checkpoint 19 (85.25%), dropped to 84.02% after retraining, and later reached 86.07%. Fig. 4 presents the corresponding progression over time.

Model scaling: We observed diminishing returns for the 7B prover as its improvement rate on the training corpus and miniF2F began to saturate, suggesting a capacity limitation inherent to the model size. To verify the scalability of our training recipe, we extended the multi-stage expert iteration experiments to the Qwen2.5-32B model. Scaling to 32B parameters yielded immediate gains even without additional data: checkpoints 16 and 17, trained on identical corpus, achieved 82.38% and 82.79% respectively. Critically, the 32B model demonstrated superior out-of-distribution generalization on ProofNet (41.4% vs. 34.4% for 7B), indicating that increased model capacity enhances transfer to novel mathematical domains beyond the training distribution.

3.4. Further Analysis on Inference

Computational budget: We report the accumulative performance across a small grid of search hyperparameters (varying branching factors and depth rewards) with a budget of pass@8192 per configuration. We adopt this metric to rigorously estimate the system’s maximum reasoning capability. Notably, our system remains robust without this accumulation: a single fixed configuration (branching factor 3 with depth reward 2 for miniF2F, or branching factor 8 with depth reward 1 for ProofNet) yields comparable results given a budget of pass@16384. The accumulative metric thus serves to capture the long tail of complex problems that reside at the boundaries of standard search parameters, providing a comprehensive view of the model’s reachable

proof set.

Planner effectiveness: The integration of the Planner agent provides a massive performance boost at both model scales. The 7B model’s performance jumps from 82.4% to 92.6% with the planner, while the 32B model improves from 86.1% to 95.1%. Notably, the planner narrows the performance gap between model scales, demonstrating that effective search strategies can partially compensate for raw model capacity.

Subgoal cache: The shared subgoal cache is critical for reducing computational complexity. It effectively transforms the search complexity from the product of individual subgoal search spaces to their sum. Empirically, for `amc12a.2008_p25` and `mathd_algebra.17` in miniF2F-test, a planner-only setup without caching failed to find a solution after 8,192 cumulative prover instances. In contrast, with the shared subgoal cache, the system consistently solved these problems using fewer than 512 cumulative instances, preventing redundant computation on established subgoals.

4. Conclusion

The primary contributions of this work are the design, implementation, and empirical validation of a holistic system for scaling LLM-based step-provers. On the training side, our multi-stage expert iteration pipeline overcomes common performance plateaus and enables sustained improvement over an extended training period. On the inference side, by leveraging a planner agent for subgoal decomposition and a shared subgoal cache for collaborative search, our system transforms intractable search spaces into manageable sequences of tasks. Empirically, `BFS-Prover-V2` not only achieves state-of-the-art performance on miniF2F but also exhibits robust OOD generalization on ProofNet, providing strong evidence for the efficacy of our approach.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Anthony, T., Tian, Z., and Barber, D. Thinking fast and slow with deep learning and tree search. *Advances in neural information processing systems*, 30, 2017.

Azerbaiyev, Z., Piotrowski, B., Schoelkopf, H., Ayers, E. W., Radev, D., and Avigad, J. Proofnet: Autoformalizing and formally proving undergraduate-level mathematics. *arXiv preprint arXiv:2302.12433*, 2023.

Baba, K., Liu, C., Kurita, S., and Sannai, A. Prover agent: An agent-based framework for formal mathematical proofs. *arXiv preprint arXiv:2506.19923*, 2025.

Bibel, W., Hölldobler, S., and Neugebauer, G. *Deduction: automated logic*. Academic Press London, 1993.

Blokpoel, M. mathlib: A scala package for readable, verifiable and sustainable simulations of formal theory. *Journal of Open Source Software*, 9(99):6049, 2024.

Cao, C., Song, L., Li, Z., Le, X., Zhang, X., Xue, H., and Yang, F. Reviving dsp for advanced theorem proving in the era of reasoning models. *arXiv preprint arXiv:2506.11487*, 2025.

Chen, J., Chen, W., Du, J., Hu, J., Jiang, Z., Jie, A., Jin, X., Jin, X., Li, C., Shi, W., et al. Seed-prover 1.5: Mastering undergraduate-level theorem proving via learning from experience. *arXiv preprint arXiv:2512.17260*, 2025a.

Chen, L., Gu, J., Huang, L., Huang, W., Jiang, Z., Jie, A., Jin, X., Jin, X., Li, C., Ma, K., et al. Seed-prover: Deep and broad reasoning for automated theorem proving. *arXiv preprint arXiv:2507.23726*, 2025b.

Comanici, G., Bieber, E., Schaekermann, M., Pasupat, I., Sachdeva, N., Dhillon, I., Blistein, M., Ram, O., Zhang, D., Rosen, E., et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.

Geuvers, H. Proof assistants: History, ideas and future. *Sadhana*, 34(1):3–25, 2009.

Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.

Han, J. M., Rute, J., Wu, Y., Ayers, E. W., and Polu, S. Proof artifact co-training for theorem proving with language models. *arXiv preprint arXiv:2102.06203*, 2021.

Jiang, A. Q., Welleck, S., Zhou, J. P., Li, W., Liu, J., Jamnik, M., Lacroix, T., Wu, Y., and Lample, G. Draft, sketch, and prove: Guiding formal theorem provers with informal proofs. *arXiv preprint arXiv:2210.12283*, 2022.

Lample, G., Lacroix, T., et al. Hypertree proof search for neural theorem proving. *Advances in Neural Information Processing Systems*, 35:26337–26349, 2022.

Li, J., Beeching, E., Tunstall, L., Lipkin, B., Soletskyi, R., Huang, S., Rasul, K., Yu, L., Jiang, A. Q., Shen, Z., et al. Numinamath: The largest public dataset in ai4maths with 860k pairs of competition math problems and solutions. *Hugging Face repository*, 13(9):9, 2024.

Liang, Z., Song, L., Li, Y., Yang, T., Zhang, F., Mi, H., and Yu, D. Mps-prover: Advancing stepwise theorem proving by multi-perspective search and data curation. *arXiv preprint arXiv:2505.10962*, 2025.

Lin, Y., Tang, S., Lyu, B., Wu, J., Lin, H., Yang, K., Li, J., Xia, M., Chen, D., Arora, S., et al. Goedel-prover: A frontier model for open-source automated theorem proving. *arXiv preprint arXiv:2502.07640*, 2025a.

Lin, Y., Tang, S., Lyu, B., Yang, Z., Chung, J.-H., Zhao, H., Jiang, L., Geng, Y., Ge, J., Sun, J., et al. Goedel-prover-v2: Scaling formal theorem proving with scaffolded data synthesis and self-correction. *arXiv preprint arXiv:2508.03613*, 2025b.

Liu, M., Diao, S., Lu, X., Hu, J., Dong, X., Choi, Y., Kautz, J., and Dong, Y. Prorl: Prolonged reinforcement learning expands reasoning boundaries in large language models. *arXiv preprint arXiv:2505.24864*, 2025.

Moura, L. d. and Ullrich, S. The lean 4 theorem prover and programming language. In *International Conference on Automated Deduction (CADE)*, pp. 625–635. Springer, 2021.

OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Polu, S., Han, J. M., Zheng, K., et al. Formal mathematics statement curriculum learning. *arXiv preprint arXiv:2202.01344*, 2022.

Puterman, M. L. Markov decision processes. *Handbooks in operations research and management science*, 2:331–434, 1990.

Ren, Z., Shao, Z., Song, J., Xin, H., Wang, H., Zhao, W., Zhang, L., Fu, Z., Zhu, Q., Yang, D., et al. Deepseek-prover-v2: Advancing formal mathematical reasoning via

- 495 reinforcement learning for subgoal decomposition. *arXiv*
 496 *preprint arXiv:2504.21801*, 2025.
- 497
- 498 Seed, B., Chen, J., Fan, T., Liu, X., Liu, L., Lin, Z., Wang,
 499 M., Wang, C., Wei, X., Xu, W., et al. Seed1. 5-thinking:
 500 Advancing superb reasoning models with reinforcement
 501 learning. *arXiv preprint arXiv:2504.13914*, 2025.
- 502
- 503 Silver, D. and Sutton, R. S. Welcome to the era of experi-
 504 ence. *Google AI*, 1, 2025.
- 505
- 506 Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai,
 507 M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graep-
 508 pel, T., et al. A general reinforcement learning algorithm
 509 that masters chess, shogi, and go through self-play. *Sci-*
 510 *ence*, 362(6419):1140–1144, 2018.
- 511
- 512 Song, P., Yang, K., and Anandkumar, A. Towards large
 513 language models as copilots for theorem proving in lean.
 514 *arXiv preprint arXiv:2404.12534*, 2024.
- 515
- 516 Sutton, R. S. Reinforcement learning: An introduction. A
 517 *Bradford Book*, 2018.
- 518
- 519 Team, K., Du, A., Gao, B., Xing, B., Jiang, C., Chen, C.,
 520 Li, C., Xiao, C., Du, C., Liao, C., et al. Kimi k1. 5:
 521 Scaling reinforcement learning with llms. *arXiv preprint*
 522 *arXiv:2501.12599*, 2025.
- 523
- 524 Tian, Y., Huang, R., Wang, X., Ma, J., Huang, Z., Luo, Z.,
 525 Lin, H., Zheng, D., and Du, L. Evolver: Advanc-
 526 ing automated theorem proving by evolving formalized
 527 problems via symmetry and difficulty. *arXiv preprint*
 528 *arXiv:2510.00732*, 2025.
- 529
- 530 Wang, H., Unsal, M., Lin, X., Baksys, M., Liu, J., Santos,
 531 M. D., Sung, F., Vinyes, M., Ying, Z., Zhu, Z., et al.
 532 Kimina-prover preview: Towards large formal reason-
 533 ing models with reinforcement learning. *arXiv preprint*
 534 *arXiv:2504.11354*, 2025.
- 535
- 536 Wang, R., Zhang, J., Jia, Y., Pan, R., Diao, S., Pi, R., and
 537 Zhang, T. Theoremllama: Transforming general-purpose
 538 llms into lean4 experts. *arXiv preprint arXiv:2407.03203*,
 539 2024.
- 540
- 541 Welleck, S. and Saha, R. Llmstep: Llm proofstep sugges-
 542 tions in lean. *arXiv preprint arXiv:2310.18457*, 2023.
- 543
- 544 Wu, Z., Huang, S., Zhou, Z., Ying, H., Wang, J., Lin, D., and
 545 Chen, K. Internlm2. 5-stepprover: Advancing automated
 546 theorem proving via expert iteration on large-scale lean
 547 problems. *arXiv preprint arXiv:2410.15700*, 2024a.
- 548
- 549 Wu, Z., Wang, J., Lin, D., and Chen, K. Lean-github:
 Compiling github lean repositories for a versatile lean
 prover. *arXiv preprint arXiv:2407.17227*, 2024b.
- Xin, H., Guo, D., Shao, Z., Ren, Z., Zhu, Q., Liu, B., Ruan,
 C., Li, W., and Liang, X. Deepseek-prover: Advancing
 theorem proving in llms through large-scale synthetic
 data. *arXiv preprint arXiv:2405.14333*, 2024a.
- Xin, H., Ren, Z., Song, J., Shao, Z., Zhao, W., Wang, H., Liu,
 B., Zhang, L., Lu, X., Du, Q., et al. Deepseek-prover-v1.
 5: Harnessing proof assistant feedback for reinforcement
 learning and monte-carlo tree search. *arXiv preprint*
arXiv:2408.08152, 2024b.
- Xin, R., Xi, C., Yang, J., Chen, F., Wu, H., Xiao, X., Sun,
 Y., Zheng, S., and Shen, K. Bfs-prover: Scalable best-
 first tree search for llm-based automatic theorem proving.
arXiv preprint arXiv:2502.03438, 2025.
- Yang, A., Zhang, B., Hui, B., Gao, B., Yu, B., Li, C., Liu,
 D., Tu, J., Zhou, J., Lin, J., et al. Qwen2. 5-math techni-
 cal report: Toward mathematical expert model via self-
 improvement. *arXiv preprint arXiv:2409.12122*, 2024a.
- Yang, K., Poesia, G., He, J., Li, W., Lauter, K., Chaudhuri,
 S., and Song, D. Formal mathematical reasoning: A new
 frontier in ai. *arXiv preprint arXiv:2412.16075*, 2024b.
- Yang, K., Swope, A., Gu, A., Chalamala, R., Song, P.,
 Yu, S., Godil, S., Prenger, R. J., and Anandkumar, A.
 Leandojo: Theorem proving with retrieval-augmented
 language models. *Advances in Neural Information Pro-*
cessing Systems, 36, 2024c.
- Ying, H., Wu, Z., Geng, Y., Wang, J., Lin, D., and Chen,
 K. Lean workbook: A large-scale lean problem set for-
 malized from natural language math problems. *arXiv*
preprint arXiv:2406.03847, 2024.
- Yu, Q., Zhang, Z., Zhu, R., Yuan, Y., Zuo, X., Yue, Y., Dai,
 W., Fan, T., Liu, G., Liu, L., et al. Dapo: An open-source
 llm reinforcement learning system at scale. *arXiv preprint*
arXiv:2503.14476, 2025.
- Yue, Y., Yuan, Y., Yu, Q., Zuo, X., Zhu, R., Xu, W., Chen,
 J., Wang, C., Fan, T., Du, Z., et al. Vapo: Efficient and
 reliable reinforcement learning for advanced reasoning
 tasks. *arXiv preprint arXiv:2504.05118*, 2025.
- Zheng, K., Han, J. M., and Polu, S. Minif2f: a cross-system
 benchmark for formal olympiad-level mathematics. *arXiv*
preprint arXiv:2109.00110, 2021.
- Zhou, Y., Zhao, J., Zhang, Y., Wang, B., Wang, S., Chen,
 L., Wang, J., Chen, H., Jie, A., Zhang, X., et al. Solving
 formal math problems by decomposition and iterative
 reflection. *arXiv preprint arXiv:2507.15225*, 2025.

A. Case Studies

A.1. Proof Conciseness and Tactic Proficiency

A primary advantage of our step-level proof approach over the whole-proof paradigm is a dramatic reduction in proof length, which arises from the interactive nature of our method. By engaging with the Lean environment step by step, our model captures and leverages fine-grained tactic state information. This iterative feedback loop significantly improves its ability to employ powerful, high-level tactics such as `simp`, `linarith`, and `nlinarith`, enabling it to solve complex goals in a single step rather than through lengthy, explicit derivations.

To illustrate this contrast, we present a comparative analysis of proofs for two problems, `aime_1984_p7` and `amc12_2000_p1`, generated by `BFS-Prover-V2` and `DeepSeek-Prover-V2`.

`aime_1984_p7`

For the problem `aime_1984_p7`, `BFS-Prover-V2` produces a remarkably concise proof, where a single tactic `simp [h1, h0]` effectively replaces over 2000 lines that appear in numerous `have` blocks within `DeepSeek-Prover-V2`'s proof.

BFS-Prover-V2 (step-level proof):

```
theorem aime_1984_p7
  (f : ℤ → ℤ)
  (h0 : ∀ n, 1000 ≤ n → f n = n - 3)
  (h1 : ∀ n, n < 1000 → f n = f (f (n + 5))) :
  f 84 = 997 := by
  simp [h1, h0]
```

DeepSeek-Prover-V2 (whole-proof):

```
theorem aime_1984_p7 (f : ℤ → ℤ) (h0 : ∀ n, 1000 ≤ n → f n = n - 3)
  (h1 : ∀ n, n < 1000 → f n = f (f (n + 5))) : f 84 = 997 := by
  have h2 : f 1004 = 1001 := by
    have h21 : f 1004 = 1004 - 3 := by
      apply h0
      <.> norm_num
    rw [h21]
    <.> norm_num

  have h3 : f 1003 = 1000 := by
    have h31 : f 1003 = 1003 - 3 := by
      apply h0
      <.> norm_num
    rw [h31]
    <.> norm_num

  have h4 : f 1002 = 999 := by
    have h41 : f 1002 = 1002 - 3 := by
      apply h0
      <.> norm_num
    rw [h41]
    <.> norm_num

  -- ... (approximately 2000 lines of similar calculations omitted) ...

  have h193 : f 89 = 998 := by
    have h1931 : f 89 = f (f (89 + 5)) := by
      apply h1
```

```

605     <;> norm_num
606     rw [h1931]
607     have h1932 : f (89 + 5) = f 94 := by norm_num
608     rw [h1932]
609     rw [h192]
610     <;> norm_num
611     <;> linarith

612     have h194 : f 84 = 997 := by
613       have h1941 : f 84 = f (f (84 + 5)) := by
614         apply h1
615         <;> norm_num
616         rw [h1941]
617         have h1942 : f (84 + 5) = f 89 := by norm_num
618         rw [h1942]
619         rw [h193]
620         <;> norm_num
621         <;> linarith

622     rw [h194]
623     <;> norm_num

```

amc12_2000_p1

For the problem amc12_2000_p1, BFS-Prover-V2 finishes the proof with a concise application of nlinarith, replacing the roughly 80 lines of explicit case analysis over the unknowns i , m , and o required in DeepSeek-Prover-V2's proof.

BFS-Prover-V2 (step-level proof):

```

634 theorem amc12_2000_p1
635   (i m o : ℕ)
636   (h0 : i ≠ m ∧ m ≠ o ∧ o ≠ i)
637   (h1 : i*m*o = 2001) :
638   i+m+o ≤ 671 := by
639     obtain ⟨hi, hm, h2⟩ := h0
640     obtain _ | _ | _ | i := i <;> obtain _ | _ | _ | m := m <;> obtain _ | _ | _ | o :=
641       o <;> try omega
642     any_goals
643       ring_nf at h1 ⊢; nlinarith;
644     ring_nf at * <;> try omega
645     nlinarith [mul_nonneg (by positivity : 0 ≤ i) (by positivity : 0 ≤ o), mul_nonneg (
646       by positivity : 0 ≤ m) (by positivity : 0 ≤ i), mul_nonneg (by positivity : 0 ≤ o)
647       (by positivity : 0 ≤ m)]

```

DeepSeek-Prover-V2 (whole-proof):

```

650 theorem amc12_2000_p1 (i m o : ℕ) (h0 : i ≠ m ∧ m ≠ o ∧ o ≠ i) (h1 : i * m * o = 2001)
651   :
652   i + m + o ≤ 671 := by
653
654   -- ... (approximately 80 lines of properties related to unknowns i, m, and o
655   omitted) ...
656
657   have h_main : i + m + o ≤ 671 := by
658     have h8 : i = 1 ∨ i = 3 ∨ i = 23 ∨ i = 29 ∨ i = 69 ∨ i = 87 ∨ i = 667 ∨ i = 2001
659

```

```

660 := by
661   have h81 : i | 2001 := h2
662   have h82 : i ≤ 2001 := h5
663   have h83 : i > 0 := h_pos_i
664   have h84 : i = 1 ∨ i = 3 ∨ i = 23 ∨ i = 29 ∨ i = 69 ∨ i = 87 ∨ i = 667 ∨ i =
2001 := by
665     have h85 : i | 2001 := h81
666     have h86 : i ≤ 2001 := h82
667     interval_cases i <;> norm_num at h85 ⊢ <;> omega
668     exact h84
669   have h9 : m = 1 ∨ m = 3 ∨ m = 23 ∨ m = 29 ∨ m = 69 ∨ m = 87 ∨ m = 667 ∨ m = 2001
:= by
670     have h91 : m | 2001 := h3
671     have h92 : m ≤ 2001 := h6
672     have h93 : m > 0 := h_pos_m
673     have h94 : m = 1 ∨ m = 3 ∨ m = 23 ∨ m = 29 ∨ m = 69 ∨ m = 87 ∨ m = 667 ∨ m =
2001 := by
674       have h95 : m | 2001 := h91
675       have h96 : m ≤ 2001 := h92
676       interval_cases m <;> norm_num at h95 ⊢ <;> omega
677       exact h94
678   have h10 : o = 1 ∨ o = 3 ∨ o = 23 ∨ o = 29 ∨ o = 69 ∨ o = 87 ∨ o = 667 ∨ o = 2001
:= by
679     have h101 : o | 2001 := h4
680     have h102 : o ≤ 2001 := h7
681     have h103 : o > 0 := h_pos_o
682     have h104 : o = 1 ∨ o = 3 ∨ o = 23 ∨ o = 29 ∨ o = 69 ∨ o = 87 ∨ o = 667 ∨ o =
2001 := by
683       have h105 : o | 2001 := h101
684       have h106 : o ≤ 2001 := h102
685       interval_cases o <;> norm_num at h105 ⊢ <;> omega
686       exact h104
687   rcases h8 with (rfl | rfl | rfl | rfl | rfl | rfl | rfl | rfl) <;>
688   rcases h9 with (rfl | rfl | rfl | rfl | rfl | rfl | rfl | rfl) <;>
689   rcases h10 with (rfl | rfl | rfl | rfl | rfl | rfl | rfl | rfl) <;>
690   norm_num [mul_assoc, mul_comm, mul_left_comm] at h1 h0 ⊢ <;>
691   (try omega) <;>
692   (try
693     {
694       norm_num at h0 ⊢ <;>
695       (try omega) <;>
696       (try
697         {
698           ring_nf at h1 ⊢ <;>
699           omega
700         })
701       }) <;>
702   (try
703     {
704       norm_num at h0 ⊢ <;>
705       (try omega) <;>
706       (try
707         {
708           ring_nf at h1 ⊢ <;>
709           omega
710         })
711       }) <;>
712   (try
713     {
714       norm_num at h0 ⊢ <;>
715       (try omega) <;>
716       (try
717         {

```

```

715         ring_nf at h1 ⊢ <;>
716         omega
717     })
718   }) <;>
719   (try
720     {
721       norm_num at h0 ⊢ <;>
722       (try omega) <;>
723       (try
724         {
725           ring_nf at h1 ⊢ <;>
726           omega
727         })
728     })
729   <;>
730   (try omega)
731   <;>
732   (try
733     {
734       ring_nf at h1 ⊢ <;>
735       omega
736     })
737   <;>
738   exact h_main

```

A.2. Novel Proof Strategies

Another significant advantage of our step-level proof approach is its ability to discover novel proof strategies that whole-proof or human-proof methods typically would not consider. By exploring the proof space progressively, our system can identify non-obvious connections and construct solutions that are both elegant and insightful.

We illustrate this capability by examining the problems `imo_1963_p5` and `algebra_amgm_sumltoneqn-prod1tonleq1`, each of which highlights a distinct advantage of our approach.

`imo_1963_p5`

For the problem `imo_1963_p5`, our model, DeepSeek-Prover-V2, and Compfiles dataset provide step-level proof, whole-proof, and human-proof versions, respectively. Notably, both whole-proof and human-proof approaches employ similar strategies: multiplying both sides of the equation by $2 \cdot \sin(\pi/7)$, then applying sum-to-product trigonometric identities for simplification. In contrast, BFS-Prover-V2 develops an entirely different approach: first transforming the left side of the equation into a polynomial in $\cos(\pi/7)$ using double and triple angle formulas, then proving that $\cos(\pi/7)$ satisfies the corresponding polynomial equation.

BFS-Prover-V2 (step-level proof):

```

758 theorem imo_1963_p5 :
759   Real.cos (π / 7) - Real.cos (2 * π / 7) + Real.cos (3 * π / 7) = 1 / 2 := by
760   have x : Real.pi / 7 = Real.pi / 7 * 1 := by ring
761   have h : 3 * Real.pi / 7 = Real.pi - 4 * Real.pi / 7 := by ring
762   rw [h, cos_sub] <;> norm_num
763   have h2 := cos_two_mul (Real.pi / 7)
764   have h3 := cos_three_mul (π / 7)
765   rw [show 4 * Real.pi / 7 = Real.pi - 3 * Real.pi / 7 by ring,
766     cos_sub]
767   simp [h2, h3, cos_two_mul, sin_pi, cos_pi]
768   ring_nf at h2 h3 ⊢
769   norm_num [h2, h3, cos_pi_div_two]

```

```

770 ring_nf
771   <|> have h4 := cos_pi
772   <|> simp [h4]
773 ring_nf at * <|> norm_num
774 rw [← sub_eq_zero]
775 nth_rewrite 1 [← sub_eq_zero]
776 ring_nf
777 apply eq_of_sub_eq_zero
778 let y := cos (Real.pi * (1 / 7))
779 have := cos_three_mul (Real.pi * (1 / 7))
780 ring_nf at *
781 apply eq_of_sub_eq_zero
782 clear this h3 h2
783 apply eq_of_sub_eq_zero
784 have := cos_three_mul (Real.pi * (1 / 7))
785 field_simp [mul_assoc] at *
786 on_goal 1 => ring
787 replace : Real.pi * (1 / 7 : ℝ) = Real.pi / 7 := by ring
788 try rw [this]; norm_num
789 have h5 := cos_three_mul (Real.pi / 7)
790 have : 3 * (Real.pi / 7) = Real.pi - 4 * (Real.pi / 7) := by ring
791 simp [this, cos_pi] at h5
792 let z := cos (Real.pi / 7)
793 rcases lt_trichotomy 0 z with hz | hz | hz
794 any_goals simp_all [show cos (4 * (Real.pi / 7)) = cos (2 * (2 * (Real.pi / 7))) by
795   ring,
796   cos_two_mul]
797 any_goals nlinarith [cos_sq_add_sin_sq (Real.pi / 7), Real.sin_pi_div_two_sub,
798   pow_two_nonneg (cos (Real.pi / 7) - 1), pow_two_nonneg (cos (Real.pi / 7) + 1)]
799 contrapose hz
800 refine not_lt.2 ?_
801 apply cos_nonneg_of_mem_Icc <|> constructor <|> linarith [pi_pos]

```

DeepSeek-Prover-V2 (whole-proof):

```

801 theorem imo_1963_p5 : Real.cos (Real.pi / 7) - Real.cos (2 * Real.pi / 7) + Real.cos
802   (3 * Real.pi / 7) = 1 / 2 := by
803   have h1 : 2 * Real.sin (Real.pi / 7) * (Real.cos (Real.pi / 7) - Real.cos (2 * Real.
804     pi / 7) + Real.cos (3 * Real.pi / 7)) = Real.sin (Real.pi / 7) := by
805     have h1_1 : 2 * Real.sin (Real.pi / 7) * Real.cos (Real.pi / 7) = Real.sin (2 *
806       Real.pi / 7) := by
807         have h1_1_1 : Real.sin (2 * Real.pi / 7) = 2 * Real.sin (Real.pi / 7) * Real.cos (
808           Real.pi / 7) := by
809           have h1_1_2 : Real.sin (2 * Real.pi / 7) = Real.sin (2 * (Real.pi / 7)) := by
810             ring
811             rw [h1_1_2]
812             have h1_1_3 : Real.sin (2 * (Real.pi / 7)) = 2 * Real.sin (Real.pi / 7) * Real.
813               cos (Real.pi / 7) := by
814                 rw [Real.sin_two_mul]
815                 <|> ring
816                 rw [h1_1_3]
817                 <|> ring
818             linarith
819             have h1_2 : 2 * Real.sin (Real.pi / 7) * Real.cos (2 * Real.pi / 7) = Real.sin (3 *
820               Real.pi / 7) - Real.sin (Real.pi / 7) := by
821               -- ... (approximately 20 lines of calculations omitted) ...
822               have h1_3 : 2 * Real.sin (Real.pi / 7) * Real.cos (3 * Real.pi / 7) = Real.sin (4 *
823                 Real.pi / 7) - Real.sin (2 * Real.pi / 7) := by

```

```

825 -- ... (approximately 20 lines of similar calculations omitted) ...
826
827   have h14 : Real.sin (4 * Real.pi / 7) = Real.sin (3 * Real.pi / 7) := by
828
829 -- ... (approximately 20 lines of similar calculations omitted) ...
830
831 have h2 : Real.sin (Real.pi / 7) > 0 := by
832   apply Real.sin_pos_of_pos_of_lt_pi
833   · linarith [Real.pi_pos, Real.pi_gt_three]
834   · linarith [Real.pi_pos, Real.pi_gt_three]
835
836 have h3 : Real.cos (Real.pi / 7) - Real.cos (2 * Real.pi / 7) + Real.cos (3 * Real.
837 pi / 7) = 1 / 2 := by
838   have h31 : 2 * Real.sin (Real.pi / 7) > 0 := by linarith
839   have h32 : Real.cos (Real.pi / 7) - Real.cos (2 * Real.pi / 7) + Real.cos (3 *
840 Real.pi / 7) = 1 / 2 := by
841     apply mul_left_cancel0 (show (2 * Real.sin (Real.pi / 7) : ℝ) ≠ 0 by linarith)
842     nlinarith [Real.sin_le_one (Real.pi / 7), Real.sin_le_one (2 * Real.pi / 7),
843 Real.sin_le_one (3 * Real.pi / 7),
844 Real.sin_le_one (4 * Real.pi / 7), Real.sin_le_one (Real.pi / 7)]
845     exact h32
846
847 apply h3

```

Compfiles dataset (human-proof):

```

848 theorem imo1963_p5 :
849   Real.cos (π/7) - Real.cos (2*π/7) + Real.cos (3*π/7) = 1/2 := by
850   rw [show (2*π/7) = π - (5*π/7) by linarith]
851   rw [Real.cos_pi_sub]
852   simp only [sub_neg_eq_add]
853   have h : 2 * Real.sin (π / 7) ≠ 0 := by
854     simp only [ne_eq, mul_eq_zero, OfNat.ofNat_ne_zero, false_or]
855     apply ne_of_gt
856     apply Real.sin_pos_of_pos_of_lt_pi
857     simp only [Nat.ofNat_pos, div_pos_iff_of_pos_right, Real.pi_pos]
858     trans 1
859     · rw [div_lt_one (by linarith only)]
860     · linarith only [Real.pi_le_four]
861     · linarith only [Real.pi_gt_three]
862   apply (mul_right_inj' h).mp
863   rw [left_distrib, left_distrib]
864   have prod_sum : ∀ (x y : ℝ),
865     2 * Real.sin x * Real.cos y = Real.sin (x + y) - Real.sin (y - x) := by
866     intro x y
867     rw [Real.sin_add, Real.sin_sub]
868     linarith only
869   rw [prod_sum, prod_sum, prod_sum]
870   rw [show (π / 7 + π / 7) = 2 * π / 7 by linarith only]
871   rw [show (π / 7 - π / 7) = 0 by linarith only]
872   rw [show (π / 7 + 5 * π / 7) = 6 * π / 7 by linarith only]
873   rw [show (5 * π / 7 - π / 7) = 4 * π / 7 by linarith only]
874   rw [show (π / 7 + 3 * π / 7) = 4 * π / 7 by linarith only]
875   rw [show (3 * π / 7 - π / 7) = 2 * π / 7 by linarith only]
876   rw [Real.sin_zero]
877   ring_nf
878   rw [← Real.sin_pi_sub]
879   rw [show (π - π * (6 / 7)) = π / 7 by linarith]
880   congr
881   linarith

```

algebra_amgm_sumltoneqn_prodltonleq1

For the problem `algebra_amgm_sumltoneqn_prodltonleq1`, the whole-proof model `DeepSeek-Prover-V2` adopts a standard, first-principles approach: it proceeds by manually handling cases ($n = 0$, some $a_i = 0$, all $a_i > 0$), and then takes the logarithm of the product and then applies the well-known inequality $\ln(x) \leq x - 1$ to each term, resulting in a verbose proof. In contrast, `BFS-Prover-V2` recognizes the problem as a special case of the Arithmetic Mean-Geometric Mean (AM-GM) inequality. It directly invokes the corresponding theorem from `Mathlib`, `Real.geom_mean_le_arith_mean`, demonstrating an ability to leverage high-level library theorems for a more insightful and efficient proof.

BFS-Prover-V2 (step-level proof):

```

theorem algebra_amgm_sumltoneqn_prodltonleq1
  (a : ℕ → NNReal)
  (n : ℕ)
  (h₀ : ∑ x in Finset.range n, a x = n) :
  ∏ x in Finset.range n, a x ≤ 1 := by
  have g := h₀
  revert h₀
  intro amgm
  let S := Finset.range n
  by_cases h1 : n = 0
  simp[h1]
  have hn : 0 < n := by omega
  let μ := (fun (x : ℕ) => (a x : ℝ))
  let w : ℕ → ℝ := fun _ => 1
  have w_nonneg : ∀ i, 0 ≤ w i
  simp [w, zero_le_one]
  have w_pos : 0 < ∏ i in Finset.range n, w i
  all_goals aesop (add simp w)
  have w_nonneg : ∀ x ∈ S, 0 ≤ w x
  any_goals simp [w]
  have w_pos : 0 < ∏ x ∈ S, w x
  simp [w]
  have amgm_real : (∏ x ∈ S, μ x ^ (w x : ℝ)) ^ (∑ x ∈ S, (w x : ℝ)⁻¹) ≤ ∑ x ∈ S,
    (w x * μ x : ℝ) / ∑ x ∈ S, (w x : ℝ)
  apply Real.geom_mean_le_arith_mean
  exact w_nonneg
  simp [w, S, hn]
  field_simp[μ]
  simp [μ, w, S] at amgm_real
  norm_cast at amgm_real
  simp[ amgm] at amgm_real
  by_cases H : ∏ a_1 ∈ Finset.range n, a a_1 = 0
  simp [H]
  simp [← NNReal.coe_le_coe, amgm] at amgm_real
  rw [div_self] at amgm_real
  apply le_of_not_lt
  all_goals aesop (add simp NNReal)
  rw[←NNReal.coe_one] at amgm_real
  norm_cast at amgm_real
  contrapose! amgm_real
  erw [ Real.one_lt_rpow_iff]
  left
  all_goals aesop (add simp [Finset.range])

```

DeepSeek-Prover-V2 (whole-proof):

```

935 theorem algebra_amgm_sumltonleqn_proditonleq1 (a : ℕ → NNReal) (n : ℕ)
936 (h0 :  $\sum x \text{ in } \text{Finset.range } n, a \ x = n$ ) :  $(\prod x \text{ in } \text{Finset.range } n, a \ x) \leq 1$  := by
937   have h1 :  $(\prod x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R})) \leq 1$  := by
938     have h2 :  $\sum x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R}) = n$  := by
939       exact_mod_cast h0
940     by_cases h3 : n = 0
941     simp [h3]
942     have h4 : 0 < n := Nat.pos_of_ne_zero h3
943     by_cases h5 :  $\exists i, i \in \text{Finset.range } n \wedge (a \ i : \mathbb{R}) = 0$ 
944     obtain ⟨i, hi, h6⟩ := h5
945     have h7 :  $(\prod x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R})) = 0$  := by
946       have h8 : i ∈ Finset.range n := hi
947       have h9 : (a i : ℝ) = 0 := h6
948       have h10 :  $(\prod x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R})) = 0$  := by
949         calc
950            $(\prod x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R})) = \prod x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R}) :=$ 
951             rfl
952            $= 0$  := by
953             apply Finset.prod_eq_zero h8
954             simp [h9]
955           exact h10
956     linarith
957     have h5' :  $\forall i, i \in \text{Finset.range } n \rightarrow (a \ i : \mathbb{R}) > 0$  := by
958       intro i hi
959       have h6 :  $\neg \exists i, i \in \text{Finset.range } n \wedge (a \ i : \mathbb{R}) = 0$  := h5
960       have h7 : (a i : ℝ) ≠ 0 := by
961         intro h
962         have h8 :  $\exists i, i \in \text{Finset.range } n \wedge (a \ i : \mathbb{R}) = 0$  := ⟨i, hi, by simp_all⟩
963         contradiction
964       have h8 : (a i : ℝ) > 0 := by
965         have h9 : (a i : ℝ) ≥ 0 := by exact_mod_cast (a i).prop
966         contrapose! h7
967       linarith
968     exact h8
969     have h6 :  $(\prod x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R})) \leq 1$  := by
970       have h7 :  $\sum x \text{ in } \text{Finset.range } n, \text{Real.log } ((a \ x : \mathbb{R})) \leq 0$  := by
971         have h8 :  $\sum x \text{ in } \text{Finset.range } n, \text{Real.log } ((a \ x : \mathbb{R})) \leq \sum x \text{ in } \text{Finset.}$ 
972           range n, ((a x : ℝ) - 1) := by
973             apply Finset.sum_le_sum
974             intro i hi
975             have h9 : (a i : ℝ) > 0 := h5' i hi
976             have h10 :  $\text{Real.log } ((a \ i : \mathbb{R})) \leq (a \ i : \mathbb{R}) - 1$  := by
977               have h11 :  $\text{Real.log } ((a \ i : \mathbb{R})) \leq (a \ i : \mathbb{R}) - 1$  := by
978                 linarith [Real.log_le_sub_one_of_pos h9]
979             exact h11
980           exact h10
981       -- ... (approximately 40 lines of calculations omitted) ...
982     have h2 :  $(\prod x \text{ in } \text{Finset.range } n, a \ x) \leq 1$  := by
983       have h3 :  $(\prod x \text{ in } \text{Finset.range } n, a \ x : \mathbb{R}) \leq 1$  := by
984         exact h1
985       have h4 :  $(\prod x \text{ in } \text{Finset.range } n, a \ x : \mathbb{R}) = (\prod x \text{ in } \text{Finset.range } n, a \ x : \mathbb{R}) :=$ 
986         rfl
987       have h5 :  $(\prod x \text{ in } \text{Finset.range } n, a \ x : \mathbb{R}) = (\prod x \text{ in } \text{Finset.range } n, (a \ x : \mathbb{R})) :=$ 
988         by simp
989       have h6 :  $(\prod x \text{ in } \text{Finset.range } n, a \ x : \mathbb{R}) \leq 1$  := by simp [h5] using h1
990       have h7 :  $(\prod x \text{ in } \text{Finset.range } n, a \ x : \text{NNReal}) \leq 1$  := by
991         norm_cast at h6 ⊢
992         <.> simp_all [Finset.prod_range_succ]
993         <.> norm_num
994         <.> linarith
995       simp using h7

```

```
exact h₂
```

Despite the advantages of step-level proof in significantly reducing proof length and discovering novel proof strategies, step-level proof has one notable limitation: poor readability. The interactive nature of step-level proof generation often results in proofs that are more challenging for humans to follow and understand compared to the more structured and explanatory whole-proof approaches. This trade-off between conciseness and readability represents a crucial consideration when evaluating the practical utility of different proof generation paradigms.

990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044

B. Illustration of Planner-Prover Paradigm with an IMO Problem

To demonstrate the effectiveness of our Planner-Prover paradigm, we present an analysis of the solution process for a challenging IMO problem: imo_1969_p2.

In the following proof, the statements `h_coeffs_polar`, `h_y_rewritten_with_polar`, and `h_y_collapsed_to_single_cos` represent the dynamic replanning phase, while all other have statements belong to the initial planning phase. Unlike in conventional whole-proof methods, have statements in our framework are presented without the `:= by` clause. This example highlights the crucial role of dynamic replanning in our system. Without dynamic replanning, the prover gets stuck at `h_y_is_sinusoid`, failing to complete the proof even after 7,200 attempts. With dynamic replanning, however, the system successfully completes the proof in just 800 attempts. The dynamic replanning process breaks down complex steps into smaller, more manageable subgoals, which enables the prover to bypass critical bottlenecks more efficiently.

imo_1969_p2 - Part 1

```

theorem imo_1969_p2
  (m n : ℝ)
  (k : ℕ)
  (a : ℕ → ℝ)
  (y : ℝ → ℝ)
  (h0 : 0 < k)
  (h1 : ∀ x, y x = ∑ i in Finset.range k, ((Real.cos (a i + x)) / (2^i)))
  (h2 : y m = 0)
  (h3 : y n = 0) : ∃ t : ℤ, m - n = t * Real.pi := by
  have h_cos_add : ∀ i x, Real.cos (a i + x) = Real.cos (a i) * Real.cos x - Real.sin
    (a i) * Real.sin x
  simp [cos_add, add_right_inj]

  have h_y_sum_expanded : ∀ (x : ℝ), y x = ∑ i in (Finset.range k : Finset ℕ), (Real.
    cos (a i) * Real.cos x - Real.sin (a i) * Real.sin x) / ((2 : ℕ) ^ i : ℝ)
  simp [h1, h_cos_add ]

  have h_y_sum_split : ∀ (x : ℝ), y x = ∑ i in (Finset.range k : Finset ℕ), Real.cos
    (a i) * Real.cos x / ((2 : ℕ) ^ i : ℝ) - ∑ i in (Finset.range k : Finset ℕ),
    Real.sin (a i) * Real.sin x / ((2 : ℕ) ^ i : ℝ)
  intro z <=> simp_rw [h_y_sum_expanded]
  simp [sub_div, Finset.sum_sub_distrib]

  have h_y_expand : ∀ (x : ℝ), y x = ∑ i in (Finset.range k : Finset ℕ), Real.cos (a
    i) / ((2 : ℕ) ^ i : ℝ) * Real.cos x - ∑ i in (Finset.range k : Finset ℕ), Real.
    sin (a i) / ((2 : ℕ) ^ i : ℝ) * Real.sin x
  intro x_exp
  simp only [Finset.sum_mul, h_y_sum_split]
  congr <=> symm <=> field_simp <=> ring

  have h_k_ge_one : 1 ≤ k
  apply Nat.succ_le_of_lt <=> exact h0

  have h_complex_repr : ((∑ i in (Finset.range k : Finset ℕ), Real.cos (a i) / ((2 : ℕ)
    ) ^ i : ℝ), ∑ i in (Finset.range k : Finset ℕ), Real.sin (a i) / ((2 : ℕ) ^ i : ℝ)
    ) : ℂ = ∑ i in (Finset.range k : Finset ℕ), Complex.exp (↑(a i) * Complex.I) / ↑
    (((2 : ℕ) ^ i) : ℝ)
  simp [Complex.exp_mul_I, div_eq_inv_mul, Complex.ext_iff]
  simp [Complex.cos_ofReal_re, Complex.sin_ofReal_re] <=> field_simp <=> norm_cast
  constructor <=> apply Finset.sum_congr <=> aesop
  field_simp [_root_.pow_add, show (4 : ℝ) = 2 ^ 2 by norm_num] <=> ring
  norm_num [mul_comm _ 2, pow_mul]
  rewrite [show (4 : ℝ) ^ x = (2 * 2 : ℝ) ^ x by ring, mul_pow]
  <=> field_simp
  <=> ring

```

```

1100 have h_sum_split :  $\sum i \text{ in } (\text{Finset.range } k : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R}) = \text{Complex.exp } (\uparrow(a \ 0) * \text{Complex.I}) + \sum i \text{ in } (\text{Finset.Icc } 1 \ (k-1) : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R})$ 
1101
1102
1103 have h_range_split :  $\text{Finset.range } k = \text{insert } 0 \ (\text{Finset.Icc } 1 \ (k - 1))$ 
1104
1105 ext x <|> simp [Nat.lt_succ_iff]
1106
1107 rcases x with (|_|x) <|> omega
1108
1109 rw [h_range_split, Finset.sum_insert]
1110
1111 norm_num [pow_zero, eq_self_iff_true]
1112
1113 simp [Nat.le_zero]
1114
1115
1116 have h_abs_head :  $\text{Complex.abs } (\text{Complex.exp } (\uparrow(a \ 0) * \text{Complex.I})) = 1$ 
1117
1118 simp [Complex.abs_exp, eq_self_iff_true]
1119
1120
1121 have h_tail_geom_sum_val :  $\sum i \text{ in } (\text{Finset.Icc } 1 \ (k - 1) : \text{Finset } \mathbb{N}), 1 / ((2 : \mathbb{N}) ^ i : \mathbb{R}) = 1 - 1 / (2 : \mathbb{R}) ^ (k - 1)$ 
1122
1123 have h_tight :  $(1 : \mathbb{R}) \leq k$ 
1124
1125 norm_cast at * <|>
1126
1127 linarith
1128
1129 clear h_tight h_sum_split h_complex_repr h_y_expand h_y_sum_split h_y_sum_expanded
1130
1131 h_cos_add h2 h3 h1 h0
1132
1133 induction' k <|> simp [Finset.sum_Icc_succ_top, *]
1134
1135 induction'  $\langle \mathbb{N} \rangle$  <|> simp_all [Finset.sum_Icc_succ_top, pow_succ]
1136
1137 ring
1138
1139 <|>ring_nf
1140
1141
1142 have h_abs_tail_le :  $\text{Complex.abs } \sum i \text{ in } (\text{Finset.Icc } 1 \ (k-1) : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R}) \leq 1 - 1 / (2 : \mathbb{R}) ^ (k - 1)$ 
1143
1144 rw [← h_tail_geom_sum_val]
1145
1146 apply (Complex.abs.sum_le _ _).trans_eq
1147
1148 apply Finset.sum_congr rfl
1149
1150 intro i _
1151
1152 simp [Complex.abs_exp_ofReal_mul_I, Nat.cast_pow, Nat.cast_ofNat]
1153
1154
1155 have h_abs_tail_lt_one :  $\text{Complex.abs } \sum i \text{ in } (\text{Finset.Icc } 1 \ (k-1) : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R}) < 1$ 
1156
1157 refine lt_of_le_of_lt h_abs_tail_le ?_
1158
1159 refine sub_lt_self _ (by positivity)
1160
1161
1162 have h_abs_ge_by_rev_triangle :  $\text{Complex.abs } \sum i \text{ in } (\text{Finset.range } k : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R}) \geq 1 - \text{Complex.abs } \sum i \text{ in } (\text{Finset.Icc } 1 \ (k-1) : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R})$ 
1163
1164 rw [h_sum_split]
1165
1166 rw [← h_abs_head]
1167
1168 apply Complex.abs.le_add

```

imo_1969_p2 - Part 2

```

1144 have h_abs_ge_final :  $\text{Complex.abs } \sum i \text{ in } (\text{Finset.range } k : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R}) \geq 1 / (2 : \mathbb{R}) ^ (k-1)$ 
1145
1146 refine' _root_.trans h_abs_ge_by_rev_triangle _
1147
1148 linarith [h_abs_tail_le]
1149
1150
1151 have h_abs_gt_zero :  $0 < \text{Complex.abs } \sum i \text{ in } (\text{Finset.range } k : \text{Finset } \mathbb{N}), \text{Complex.exp } (\uparrow(a \ i) * \text{Complex.I}) / \uparrow(((2 : \mathbb{N}) ^ i) : \mathbb{R})$ 
1152
1153 linarith [pow_two_nonneg ((k - 1 :  $\mathbb{N}$ ) :  $\mathbb{R}$ )]
1154
1155
1156 have h_complex_val_ne_zero :  $(\sum i \text{ in } (\text{Finset.range } k : \text{Finset } \mathbb{N}), \text{Real.cos } (a \ i) /$ 

```

```

1155 ((2 : ℕ) ^ i : ℝ), ∑ i in (Finset.range k : Finset ℕ), Real.sin (a i) / ((2 : ℕ) ^
1156 i : ℝ) : ℂ ≠ 0
1157 focus all_goals (norm_num; aesop)

```

```

1160 have h_coeffs_polar : ∃ (R b : ℝ), 0 < R ∧ (∑ i in (Finset.range k : Finset ℕ),
1161   Real.cos (a i) / ((2 : ℕ) ^ i : ℝ)) = R * Real.cos b ∧ (∑ i in (Finset.range
1162   k : Finset ℕ), Real.sin (a i) / ((2 : ℕ) ^ i : ℝ)) = R * Real.sin b
1163 set x := ∑ i ∈ Finset.range k, cos (a i) / ((2 : ℝ) ^ i)
1164 use Complex.abs (∑ i ∈ Finset.range k, Complex.exp (↑(a i) * Complex.I) / ↑(2
1165   ^ i))
1166 let y : ℝ := ∑ i ∈ Finset.range k, sin (a i) / 2^i
1167 have h := Complex.abs_mul_cos_add_sin_mul_I (∑ i in Finset.range k, Complex.exp
1168   ((a i : ℝ) * Complex.I) / (2 : ℂ) ^ i)
1169 use Complex.arg (∑ i in Finset.range k, Complex.exp (↑(a i) * Complex.I) / (2 :
1170   ℝ) ^ i)
1171 simp_all [Complex.ext_iff]

```

```

1174 have h_y_rewritten_with_polar : ∃ (R a : ℝ), 0 < R ∧ ∀ x, y x = R * Real.cos a
1175   * Real.cos x - R * Real.sin a * Real.sin x
1176 obtain ⟨R, phi, hR_pos, h_cos_eq1, h_sin_eq1⟩ := h_coeffs_polar
1177 use R, phi <;> simp_all [Complex.exp_mul_I, Complex.abs]

```

```

1182 have h_y_collapsed_to_single_cos : ∃ (R a : ℝ), 0 < R ∧ ∀ x, y x = R * Real.
1183   cos (x + a)
1184 rcases h_y_rewritten_with_polar with ⟨R, a', h_R_pos, h_y_⟩
1185 use R, a', h_R_pos <;> intros <;> simp [h_y_, cos_add] <;> ring

```

```

1188 have h_y_is_sinusoid : ∃ (R a : ℝ), 0 < R ∧ (∀ x, y x = R * Real.cos (x - a))
1189 obtain ⟨R, a, _, hy⟩ := h_y_collapsed_to_single_cos
1190 use R, -a <;> aesop
1191
1192 have h_roots_exist : ∃ (R a : ℝ), 0 < R ∧ y m = R * Real.cos (m - a) ∧ y n = R *
1193   Real.cos (n - a)
1194 rcases h_y_is_sinusoid with ⟨R, a, h_R_pos, h_y_R_a⟩
1195 exact ⟨R, a, h_R_pos,
1196   by simp [h_y_R_a], by simp [h_y_R_a]⟩
1197
1198 have h_cos_zero : ∃ (R a : ℝ), 0 < R ∧ Real.cos (m - a) = 0 ∧ Real.cos (n - a) = 0
1199 rcases h_roots_exist with ⟨R, a, h_rPos, h_mEq, h_nEq⟩
1200 exact
1201   ⟨R, a, h_rPos,
1202     by have := h2; have := h3; field_simp [h1] at * <;> nlinarith,
1203     by have := h3; have := h2; field_simp [h1] at * <;> nlinarith)
1204
1205 have h_roots_in_pi_half_multiples : ∃ (a : ℝ) (t1 t2 : ℤ), m - a = (2 * (t1 : ℝ) +
1206   1) * Real.pi / 2 ∧ n - a = (2 * (t2 : ℝ) + 1) * Real.pi / 2
1207 rcases h_cos_zero with ⟨R, a, _, h_m_cos_zero, h_n_cos_zero⟩
1208 rw [cos_eq_zero_iff] at h_m_cos_zero h_n_cos_zero
1209 exact ⟨a, ↑(Classical.choose h_m_cos_zero), ↑(Classical.choose h_n_cos_zero)
1210   , by convert h_m_cos_zero.choose_spec, by convert h_n_cos_zero.choose_spec⟩

```

```

1210
1211   have h_m_minus_n_form :  $\exists t_1 t_2 : \mathbb{Z}, m - n = ((2 * (t_1 : \mathbb{R}) + 1) * \text{Real.pi} / 2) -$ 
1212      $((2 * (t_2 : \mathbb{R}) + 1) * \text{Real.pi} / 2)$ 
1213   obtain  $\langle z, t_1, t_2, h\_z\_root\_m, h\_z\_root\_n \rangle := h\_roots\_in\_pi\_half\_multiples$ 
1214   refine  $\langle t_1, t_2, ?\_ \rangle <;>$ 
1215   linarith
1216
1217   have h_m_minus_n_simplified :  $\exists t_1 t_2 : \mathbb{Z}, m - n = (\uparrow(t_1 - t_2) : \mathbb{R}) * \text{Real.pi}$ 
1218   rcases h_m_minus_n_form with  $\langle t_1, t_2, h\_form \rangle <;>$ 
1219     exists t1 <;> exists t2 <;> field_simp at h_form  $\vdash <;>$  linarith
1220
1221   obtain  $\langle t_1, t_2, h\_m\_sub\_n\_t_1\_t_2 \rangle := h\_m\_minus\_n\_simplified <;>$  use t1 - t2 <;>
1222     linarith [h_m_sub_n_t1_t2]
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264

```

C. Prompts Used in This Work

C.1. Prompts for Autoformalization

Our autoformalization pipeline operates in two stages to ensure syntactic correctness. First, an `Initial Formalization Prompt` (shown below) translates a natural language problem into a Lean 4 theorem statement. If the generated code fails to compile, an `Error Feedback Prompt` is then deployed to revise the statement, using the verbatim error message from the Lean compiler as direct feedback for revision.

Prompt for Initial Formalization

You are an expert in math proof and the theorem prover: Lean. Given a math problem that contains the question and all conditions, and its corresponding solution that contains solution steps and the correct answer, generate a mathematically equivalent proof problem and rewrite it in the Lean 4 statement. You should follow the following procedures.

- a): Identify all questions and conditions in the given problem.
- b): Identify all solution steps and the correct answers in the given solution.
- c): With the questions and conditions in a) and correct answers in b), translate the (question, conditions, correct answer) tuple to a mathematically equivalent proof problem that proves `question == answer` given conditions.
- d): Rewrite the math proof problem in c) to a Lean 4 statement. Note that you should write the statement only, no proof is required. This also means you do not need to consider the solution steps either.

The first priority is to ensure the generated Lean code can be built successfully. Consider using the following tips.

- Use a broader import, e.g., `import Mathlib`, to bring in the entirety of the necessary library, and remove specific import of submodules, e.g., `import Mathlib.LinearAlgebra.BasicReal3`, accordingly.
- Add `noncomputable` before `def` only when necessary.
- Use `by` instead of `begin end`.
- Add `sorry` to skip the proof.

You should strictly follow the below criteria to guarantee the lean statement is equivalent to the mathematical problem.

- Each definition used in Lean 4 statement should only directly appear in the conditions problem in a).
- Each definition should NOT come from and assume any knowledge directly from the solution step in b).
- Each condition in a) should be used as a definition in Lean 4.
- For any implications appearing in the conclusions of the original problem, extract their antecedents and declare them as explicit assumptions before the colon, leaving only the consequent in the conclusion after the colon.
- For equations, structure the theorem in the form 'conditions : conclusions' where conditions include variable definitions and domains, and conclusions are the solutions to the equation, avoiding implication or equivalence symbols.

Below are examples to illustrate the process:

Example 1 (Number Theory):

Lean 4 statement:

```

1320 theorem nt3_problem (n p : ℕ) (hn : n > 1) (hp : Nat.Prime p)
1321   (h1 : n ∣ (p - 1)) (h2 : p ∣ (n^6 - 1)) :
1322   ∃ k : ℕ, (p - n = k^2) ∨ (p + n = k^2) := by
1323   sorry
1324

```

problem:

NT3. Let $n > 1$ be a positive integer and p a prime number such that $n \mid (p - 1)$ and $p \mid (n^6 - 1)$. Prove that at least one of the numbers $p - n$ and $p + n$ is a perfect square.

Example 2 (Number Theory):**Lean 4 statement:**

```

1331 theorem nt4_problem (x y : ℕ)
1332   (hx : x > 0) (hy : y > 0)
1333   (h1 : ∃ m : ℕ, 3 * x + 4 * y = m^2)
1334   (h2 : ∃ n : ℕ, 4 * x + 3 * y = n^2) :
1335   7 ∣ x ∧ 7 ∣ y := by
1336   sorry
1337

```

problem:

NT4. If the positive integers x and y are such that both $3x + 4y$ and $4x + 3y$ are perfect squares, prove that both x and y are multiples of 7.

Example 3 (Algebra):**Lean 4 statement:**

```

1341 theorem sum_not_zero (a b c d : ℝ)
1342   (h1 : a * b * c - d = 1)
1343   (h2 : b * c * d - a = 2)
1344   (h3 : c * d * a - b = 3)
1345   (h4 : d * a * b - c = -6) :
1346   a + b + c + d ≠ 0 := by
1347   sorry
1348

```

problem:

The real numbers a, b, c, d satisfy simultaneously the equations $abc - d = 1, bcd - a = 2, cda - b = 3, dab - c = -6$. Prove that $a + b + c + d \neq 0$.

Example 4 (Inequality):**Lean 4 statement:**

```

1351 theorem inequality_proof (a b c : ℝ)
1352   (ha : a > 0) (hb : b > 0) (hc : c > 0) :
1353   8 / ((a + b)^2 + 4*a*b*c) +
1354   8 / ((b + c)^2 + 4*a*b*c) +
1355   8 / ((c + a)^2 + 4*a*b*c) +
1356   a^2 + b^2 + c^2 ≥
1357   8 / (a + 3) + 8 / (b + 3) + 8 / (c + 3) := by
1358   sorry
1359

```

problem:

The real numbers a, b, c, d satisfy simultaneously the equations $abc - d = 1, bcd - a = 2, cda - b = 3, dab - c = -6$. Prove that $a + b + c + d \neq 0$.

Now, use the same process for the following problem and solution:

{problem}

{solution}

Prompt for Error Feedback

You are an expert in math proof and the theorem prover: Lean. You are given the following math problem that contains the question and all conditions, and its corresponding solution that contains solution steps and the correct answer.

{**problem**}

{**solution**}

A mathematically equivalent proof problem that proves question == answer given conditions is generated and rewritten in the Lean 4 statement, as shown below:

{**Lean 4 statement**}

However, this lean code got error with `lake build`, and here is the error message:

{**error message**}

Please modify the lean code to ensure it can be built successfully with `lake build`. Here is a few tips that might help:

- Use a broader import, e.g., `import Mathlib`, to bring in the entirety of the necessary library, and remove specific import of submodules, e.g., `import Mathlib.LinearAlgebra.BasicReal3`, accordingly.
- Add `noncomputable` before `def` only when necessary.
- Use `by` instead of `begin end`.
- Add `sorry` to skip the proof.

C.2. Prompts for Planner**Prompt for Initial Planning**

You are an expert assistant specializing in Math Olympiads and the Lean 4 theorem prover. Your primary goal is to generate **syntactically perfect, type-checkable** Lean 4 intermediate step code snippets (**plan**) for a given theorem. It is crucial to strictly adhere to the following rules—any violation will be considered an error.

Task

Given the following Lean 4 theorem tactic state, generate the core intermediate subgoals (`have` statements) needed for the proof.

Mandatory Rules

You must comply with every rule in this section. Failure to adhere to any single rule will result in an incorrect output.

1. Critical Rule: Explicitly Specify Set/Finset Types

This is the most common and fatal point of error. You must explicitly declare the type for any `Set` or `Finset` literal. This rule is non-negotiable.

- Correct: `{{ -1, 0, 1 }} : Set ℤ`
- Incorrect: `{ {-1, 0, 1}}`

2. Omit the Proof: Never provide the proof. Only state the `have` statement itself.**3. Valid Lean 4 Code:** The entire output block must be type-checkable in a Lean 4.10.0 environment.**4. Use Existing Names:** Use the exact, existing lemma and definition names from `Mathlib`. Do not invent names.

- 1430 5. **No Undeclared Variables:** Do not introduce any variables or constants not declared in the original theorem
1431 statement.
- 1432 6. **Explicit Multiplication:** Multiplication must always use the `*` symbol.
1433
1434 - Correct: `a * x`
1435 - Incorrect: `ax`
1436
- 1437 7. **No Chained Inequalities:** Never use chained inequalities. They must be split using logical AND `∧`.
1438
1439 - Correct: `a <= x ∧ x <= b`
1440 - Incorrect: `a <= x <= b`
1441
- 1442 8. **Correct Logarithm Function:** `Real.log` is only for the natural logarithm. For logarithms with a specified
1443 base, you must use `Real.logb`.
1444 - Correct: `Real.logb (2 : ℝ) 8`
1445 - Incorrect: `Real.log (2 : ℝ) 8`
1446
- 1447 9. **Factorial Notation:** In Lean, factorials must be written as `(n)!` or `Nat.factorial n`, not `n!`.
1448
1449 - Correct: `(n)!` or `Nat.factorial n`
1450 - Incorrect: `n!`
1451
- 1452 10. **Numeric Types Must Be Explicitly Annotated:** To avoid type ambiguity in Lean, any expression involving
1453 numeric operations must have at least one number's type specified.
1454
1455 - For division: `(1 : ℝ) / 2 = 0.5`, but `(1 : ℤ) / 2 = 0`.
1456 - For subtraction: `(1 : ℤ) - 2 = -1`, but `(1 : ℕ) - 2 = 0`.
1457 - Correct: `(a : ℝ) / b`, `a / (b : ℝ)`, `(n : ℤ) - m`
1458 - Incorrect: `a / b`, `n - m`
1459
- 1460 11. **Interval Notation:** Do not use `Icc`, `Ioo`, `Ico`, `Ioc`, etc., to represent intervals. Only use inequalities.
1461
1462 - Correct: `a <= x ∧ x <= b`
1463 - Incorrect: `Icc a b`
1464
- 1465 12. **Complex Numbers:** Use `Complex.I` for the imaginary unit and `Complex.abs` for the modulus/absolute
1466 value of a complex number.
- 1467 13. **Avoid Common Inequality Theorems:** Avoid using common inequality theorems like Holder's or Jensen's. For
1468 inequality problems, try to ensure each proof step only requires basic simplification.
- 1469 14. **Proving Equivalences:** For proofs of equivalences (iff), ensure each have statement is an implication, where
1470 the antecedent is the left side of the equivalence (when proving left-to-right) or the right side (when proving
1471 right-to-left).
1472
- 1473 15. **Real.pi Notation:** Consistently use `Real.pi` instead of `π`.
1474
- 1475 16. **Final Check:** Before providing the plan, perform a final review to ensure you have scrupulously followed all the
1476 rules above, especially the critical rule regarding `Set/Finset`.
1477

1478 **Below are examples to illustrate the process:**

1479 **Example 1:**

1480 **Theorem:**

1481
1482
1483
1484

```
1485 theorem singapore2019_r1_p7 (x : ℝ) (hx : Real.tan x = 5) :
1486   (6 + Real.sin (2 * x)) / (1 + Real.cos (2 * x)) = 83 := by
1487
```

1488 **Plan:**

```
1489
1490 have h1 : Real.sin x = 5 * Real.cos x
1491 have h2 : Real.sin x ^ 2 = 25 * Real.cos x ^ 2
1492 have h3 : 26 * Real.cos x ^ 2 = 1
1493 have hsin2x_val : Real.sin (2 * x) = (5 : ℝ) / (13 : ℝ)
1494 have hcos2x_val : Real.cos (2 * x) = -(12 : ℝ) / (13 : ℝ)
1495
```

1496 **Example 2:**

1497 **Theorem:**

```
1498 theorem problem4
1499   (g : ℕ → ℝ)
1500   (h : ∀ k : ℕ, 5 ≤ k → k ≤ 124 → g k = (Real.logb (k : ℝ) ((7 : ℝ) ^ (k ^ 2 - 1))) /
1501     (Real.logb ((k + 1 : ℝ) (7 : ℝ) ^ (k ^ 2 - 4)))) :
1502   (∏ k in Finset.Icc (5 : ℕ) 124, g k) = (41 : ℝ) / 7 := by
1503
```

1504 **Plan:**

```
1505 have h_prod_split : (∏ k in (Finset.Icc 5 124 : Finset ℕ), g k) = (∏ k in (Finset.
1506   Icc 5 124 : Finset ℕ), ((k ^ 2 - 1) / (k ^ 2 - 4 : ℝ))) * (∏ k in (Finset.Icc 5
1507   124 : Finset ℕ), (Real.logb (k : ℝ) (7 : ℝ) / Real.logb ((k + 1 : ℝ) (7 : ℝ)))
1508   have h_telescope_part1 : (∏ k in (Finset.Icc 5 124 : Finset ℕ), ((k ^ 2 - 1) / (k ^
1509   2 - 4 : ℝ))) = (41 : ℝ) / 21
1510   have h_telescope_part2 : (∏ k in (Finset.Icc 5 124 : Finset ℕ), (Real.logb (k : ℝ)
1511   (7 : ℝ) / Real.logb ((k + 1 : ℝ) (7 : ℝ))) = 3
1512   have h_final_product : (41 / 21 : ℝ) * 3 = (41 : ℝ) / 7
1513
```

1514 **Example 3:**

1515 **Theorem:**

```
1516 theorem amc12b_variant_p13
1517   (S : Finset ℝ)
1518   (h0 : ∀ (x : ℝ), x ∈ S ↔ 0 < x ∧ x ≤ 2 * Real.pi ∧ 2 - 4 * Real.sin x + 3 * Real.cos
1519     (3 * x) = 0) :
1520   S.card = 4 := by
1521
```

1522 **Plan:**

```
1523 have h_interval1 : ∃ x, 0 ≤ x ∧ x < Real.pi / 2 ∧ (2 - 4 * Real.sin x + 3 * Real.cos
1524   (3 * x) = 0)
1525 have h_interval2 : ∃ x, Real.pi / 2 ≤ x ∧ x < 3 * Real.pi / 4 ∧ (2 - 4 * Real.sin x
1526   + 3 * Real.cos (3 * x) = 0)
1527 have h_interval3 : ∃ x, 3 * Real.pi / 4 ≤ x ∧ x < Real.pi ∧ (2 - 4 * Real.sin x + 3
1528   * Real.cos (3 * x) = 0)
1529 have h_interval4 : ∃ x, Real.pi ≤ x ∧ x < 2 * Real.pi ∧ (2 - 4 * Real.sin x + 3 *
1530   Real.cos (3 * x) = 0)
1531 have h_card_eq_4 : S.card = 4
1532
```

1533 Now, use the same process for the following theorem:

1534 **{theorem}**

1535 You must follow all the mandatory rules above. After deep consideration, provide the Lean 4 intermediate step code snippets. While ensuring correctness, the more intermediate steps, the better.

Prompt for Dynamic Replanning

You are an expert assistant specializing in Math Olympiads and the Lean 4 theorem prover, with a particular talent for proof decomposition and overcoming difficult proof steps.

Your primary goal is to refine an existing proof plan by inserting more granular, logically sound subgoals to help a prover overcome a specific, identified bottleneck.

It is crucial to strictly adhere to the following rules—any violation will be considered an error.

Task

Given a Lean 4 theorem, its initial proof plan, and a specific `have` statement where a prover has become stuck, your task is to generate a new, **complete proof plan**.

This new plan must include all the original steps, but with additional, simpler `have` statements inserted **immediately before** the 'stuck' subgoal. These new steps must logically lead to the proof of the stuck subgoal, breaking down the complex reasoning into a series of more manageable steps.

Mandatory Rules

You must comply with every rule in this section. Failure to adhere to any single rule will result in an incorrect output.

(The first 16 rules are identical to those in the `Prompt for Initial Planning` and must be strictly followed.)

In addition, the following task-specific rules apply:

17. **Insert Before Stuck Step:** The new auxiliary `have` statements must be inserted **immediately before** the provided 'stuck' subgoal.
18. **Provide Complete Plan:** The output must be the **entire, updated plan**, including all original and new `have` statements in the correct order. Do not output only the new steps.
19. **Logical Progression:** The newly inserted steps must be logically sound and serve as direct prerequisites for proving the stuck subgoal. They should bridge the logical gap that caused the prover to get stuck.

Below is an example to illustrate the process:

Theorem:

```
theorem trig_identity_4x (x : ℝ) :
  Real.sin (4 * x) = 4 * Real.sin x * Real.cos x * (1 - 2 * Real.sin x ^ 2) := by
```

Initial Plan:

```
have h_sin4x_is_2sin2xcos2x : Real.sin (4 * x) = 2 * Real.sin (2 * x) * Real.cos (2 * x)
have h_final_identity : 2 * Real.sin (2 * x) * Real.cos (2 * x) = 4 * Real.sin x *
  Real.cos x * (1 - 2 * Real.sin x ^ 2)
```

Stuck Subgoal:

```
have h_final_identity : 2 * Real.sin (2 * x) * Real.cos (2 * x) = 4 * Real.sin x *
  Real.cos x * (1 - 2 * Real.sin x ^ 2)
```

Refined Plan (Your Output):

```
have h_sin4x_is_2sin2xcos2x : Real.sin (4 * x) = 2 * Real.sin (2 * x) * Real.cos (2 * x)
have h_sin2x : Real.sin (2 * x) = 2 * Real.sin x * Real.cos x
have h_cos2x_in_terms_of_sin_cos : Real.cos (2 * x) = Real.cos x ^ 2 - Real.sin x ^ 2
```

```
1595 have h_cos2x_in_terms_of_sin : Real.cos (2 * x) = 1 - 2 * Real.sin x ^ 2
1596 have h_final_identity : 2 * Real.sin (2 * x) * Real.cos (2 * x) = 4 * Real.sin x *
1597 Real.cos x * (1 - 2 * Real.sin x ^ 2)
```

Now, use the same process for the following three items:

{theorem}

{initial_plan}

{stuck_subgoal}

You must follow all the instructions and mandatory rules above. After deep consideration, provide the complete, refined Lean 4 plan.

1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649