

GAR: GENERATIVE ADVERSARIAL REINFORCEMENT LEARNING FOR FORMAL THEOREM PROVING

Ruida Wang¹, Jiarui Yao¹, Rui Pan¹, Shizhe Diao², Tong Zhang¹

¹University of Illinois Urbana-Champaign, ²NVIDIA
 {ruidaw, jiarui14, ruip4}@illinois.edu
 shizhe.diao@gmail.com
 tongzhang@tongzhang-ml.org

ABSTRACT

Solving math problems through verifiable languages such as Lean has significantly impacted both the mathematics and computer science communities. Current state-of-the-art models are often trained with expensive online Reinforcement Learning (RL) or expert iteration. However, these approaches rely on fixed problem sets, which causes inefficient training and limits the model to tackle complex problems. To overcome these limitations, we propose **GAR**: *Generative Adversarial Reinforcement learning*, a comprehensive RL training framework that jointly trains the problem composer and solver in an adversarial loop. **GAR** introduces an implicit curriculum learning mechanism, which aligns task difficulty with the prover’s evolving capability. It thereby improves the training efficiency and enables stronger performance of proving advanced theorems. Experiments show that with **GAR** training, Goedel-Prover-V2-8B and DeepSeek-Prover-V2-7B achieve an average relative improvement in pass@32 of **4.20%** on MiniF2F-Test benchmark, while DeepSeek-Prover-V2’s pass@32 on ProofNet-Test increases from 22.58% to **25.81%**. Beyond formal proving, **GAR** establishes a general RL paradigm for co-evolution of problem generation and solving under verifiable environments. The training code for this paper is open-sourced in <https://github.com/RickySkywalker/GAR-Official>

1 INTRODUCTION

The capability to perform formal mathematical reasoning has long been regarded as both a hallmark of human intelligence and a key objective of machine learning (Newell & Simon, 1956). The ability is typically assessed through rigorous mathematical derivations (Yang et al., 2024b). With the emergence of Large Language Models (LLMs), developing accurate and reliable reasoning has become an active area of research. Recent progress in ZERO RL training (Guo et al., 2025) has further advanced reasoning systems by introducing Long Chain-of-Thought (CoT) thinking models that have self-reflection and self-correction capability.

However, the inherent ambiguity of Natural Language (NL) makes it challenging to verify intermediate reasoning steps. This problem is more severe in advanced mathematics, where the task is to prove theorems rather than give a numerical or formulaic answer. The increasing complexity of modern math compounds this difficulty, as illustrated by the prolonged verification of Fermat’s Last Theorem (Wang et al., 2024). To address this issue, researchers have grounded reasoning in formal logical systems, enabling automatic verification through Formal Language (FL). Based on this idea, some researchers model the reasoning process formally with dependent type languages like Lean (De Moura et al., 2015; Moura & Ullrich, 2021) and Coq (Coq, 1996). Other uses higher-order logic to build language like Isabelle (Paulson, 1994) and HOL (Harrison, 2009). All the above languages make every intermediate step of math reasoning verifiable.

Nevertheless, mastering FLs requires considerable expertise, particularly in dependent-type systems like Lean, where proofs often demand complex type matching and the use of functions with limited data (Wang et al., 2025c). Thus, many works have sought to leverage advances in LLMs to solve FL problems and train specialized FL provers (Polu et al., 2022; Jiang et al., 2022; Xin et al., 2024a;

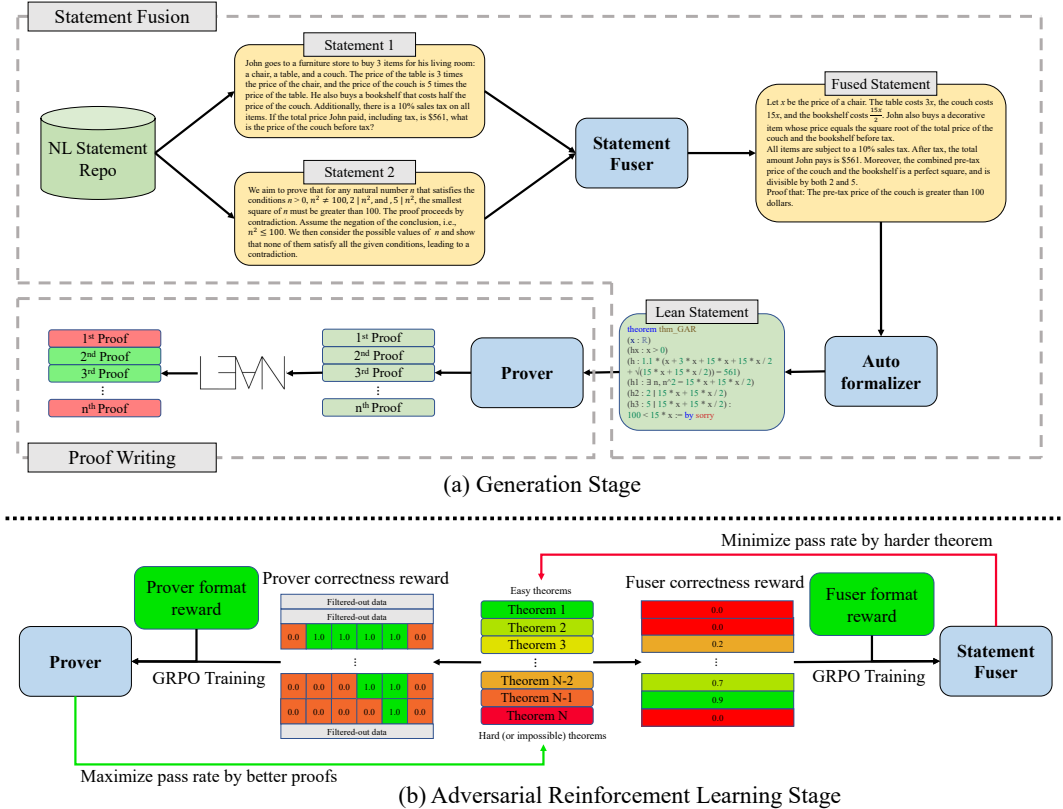


Figure 1: **GAR Training Framework:** Each iteration of **GAR** consists two stages. (a) *Generation Stage:* Pairs of NL statements are sampled from the base repository and combined by the statement fuser to create more challenging problems that fit the current model’s capability. Then, these statements are autoformalized and submitted to the prover to write multiple proofs. Subsequently, the proofs are checked by the Lean verifier for reward assignments. (b) *Adversarial Reinforcement Learning:* The prover is rewarded for producing correct proofs on medium and high-difficulty statements, while the statement fuser is rewarded for generating harder but solvable problems. This adversarial dynamic drives both models to evolve together.

Wang et al., 2024; Lin et al., 2025a; Dong & Ma, 2025b). The verifiability of FLs also motivates large-scale synthesis of new statements, which enables extensive expert iteration (Polu et al., 2022; Xin et al., 2025) or Reinforcement Learning (RL) (Ren et al., 2025; Wang et al., 2025a; Lin et al., 2025b) to further enhance prover’s performance.

However, most prior approaches in RL or expert iteration rely on fixed theorem statement sets and optimize only the prover. The statement dataset is unchanged during the process. It also leads to wasted computation on trivial or unsolvable tasks. In expert iteration, datasets often require repeated annotation. On the other hand, advanced RL methods suffer from the absence of an adaptive difficulty level of statements during rollouts, which restricts concentrated exploration and limits progress on complex theorems. More broadly, only a few works discuss a systematic way to align problem difficulty with the prover’s growing skill.

To address this limitation, we propose **GAR: Generative Adversarial Reinforcement Learning**, a comprehensive RL training framework as presented in Figure 1. **GAR** jointly optimizes the prover and the problem composer, termed as the statement fuser, through adversarial training. This process establishes an implicit curriculum learning that dynamically adjusts statement difficulty to match the prover’s development. Each iteration of the **GAR** framework consists of two stages: the generation and the adversarial RL stage. In the generation stage, the statement fuser synthesizes harder statements from existing solvable ones, and the prover attempts to solve them. In the RL training stage, the fuser is rewarded for producing difficult yet solvable statements, while the prover is re-

warded for correctly proving medium and high-difficulty problems. This adversarial setup ensures that statement complexity evolves alongside the prover’s capabilities, resulting in more efficient and effective training.

We summarize our contributions as follows: (1) We introduce **GAR**, a comprehensive RL training framework that establishes implicit curriculum learning, improving the prover’s reasoning ability while preventing wasted effort on trivial or unsolvable statements. (2) We present *Statement Fusion* technique, which enables the creation of novel formal statements beyond direct formalization of NL problems, producing theorem statements better aligned with model capability. (3) We demonstrate the effectiveness of the **GAR** through extensive experiments by training two base models using **GAR**. We achieve an average of **4.20%** relative improvement on MiniF2F-Test (Zheng et al., 2021) and raise DeepSeek-Prover-8B’s pass@32 on the more challenging ProofNet-Test (Azerbaiyev et al., 2023) from 22.58% to **25.81%**. It highlights the effectiveness and generality of **GAR**.

Furthermore, the **GAR** contributes a general RL paradigm for the co-evolution of problem generation and solving in verifiable environments, which offers new directions for adversarial co-training in other reasoning-intensive domains. To facilitate the advancement of the field, we will open-source the training and inference code of **GAR** in the near future.

2 METHODOLOGY

In this section, we introduce the **GAR** framework in detail. The prover’s task is to produce Lean4 proofs from NL-FL statement pairs. Our framework trains the prover to improve at the task by adapting theorem difficulty to the prover’s evolving capabilities through adversarial training of both the prover and the statement fuser. Each iteration of **GAR** consists of two stages: the *Generation stage* (Section 2.1) and the *Adversarial Reinforcement Learning stage* (Section 2.2). We first describe the behavior of each stage in a single iteration in detail, then present the complete procedure in Section 2.3.

2.1 GENERATION STAGE

The generation stage of **GAR** consists of two processes, namely statement fusion, which produces new statements, and proof writing, which generates and evaluates candidate proofs.

2.1.1 STATEMENT FUSION

This process generates harder statements from existing ones. It aims to mitigate the mismatch between the fixed datasets and the prover’s evolving ability. This process first samples two NL statements from the base dataset, formally: $s_{base} = (s_1^{(NL)}, s_2^{(NL)}) \sim \mathcal{D}_{stat}$, where \mathcal{D}_{stat} is a repository consists of 793,243 NL statements collect from Lean-Workbook (Wu et al., 2024b) and Numina-Math (LI et al., 2024). The sampled pair is then passed to the statement fuser trained in the previous iteration of **GAR**. The fuser is instructed to synthesize a more challenging statement by combining the key ideas of the two inputs, namely:

$$s^{(NL)} = Fuser_{i-1}(s_{base})$$

where i denotes the current iteration index, $s^{(NL)}$ is the generated NL statement, and $Fuser_{i-1}$ is the fuser from the prior step (initialized as base model when $i = 0$).

In particular, we chose thinking models like Qwen3 (Team, 2025) or DeepSeek-R1 (Guo et al., 2025) as the base model for the statement fuser because of its outstanding performance. However, direct use of their native Long CoT capability often results in overthinking and degraded quality of the generated statement. To mitigate this, we reinitialize the thinking process by passing the default thinking stage and restarting it with a dedicated indicator token named `<analysis>`. This trick yields more focused and task-specific reasoning. The full prompt for the statement fusion is shown in Figure 2. Subsequently, the generated NL statement is passed to the autoformalizer, which converts the NL statement into Lean statement $s^{(FL)}$. The formalized statement is then passed to Lean4 for compilation check. This step filters out statements with grammatical errors.

In the statement fusion stage, they deliberately separate the fusion of NL statements from the formalization of FL statements. Two key insights guide this design. First, prior studies (Wang et al., 2024;

You are an expert mathematics educator skilled in problem design. Your task is to combine multiple given problems into a SINGLE, more challenging problem by combining their key elements. Follow these steps: Please first do the following steps as your analysis process enclosed within `<analysis>`.

1. Analyze the points of knowledge that need to be used when solving the proof problem and identify overlapping or complementary aspects (e.g., shared topic areas or contrasting difficulty levels).
2. Draft the new problem that integrates at least 2 key components from each original problem and make sure the new problem requires multi-step reasoning (e.g., combining algebraic manipulation with probabilistic analysis). Also, your combined problem should have no non-trivial extension.
3. Additionally, you should make sure that the new problem is solvable.

After your analysis, you should put the new problem into an MD code block. The new problem should be a SINGLE proof problem.

Here are the statements you need to fuse:

Problem 1: `<statement_1>`

Problem 2: `<statement_2>`

Figure 2: Prompt for the Statement Fuser to generate harder statements from existing ones and restart thinking using a new indicator in **GAR** generation stage

Xin et al., 2024b) have shown that NL statements play a crucial role in supporting formal reasoning. Secondly, current 8B-scale general LLMs exhibit only a limited understanding of FL. This causes them to fail to capture essential relations, which hinders the generation of more robust statements. Consequently, the fuser produces statements that have a high likelihood of failing the compilation check if we directly fuse formal statements, as seen in Dong & Ma (2025b). Thus, **GAR** firstly fuses the statement in NL and formalizes it to FL, ensuring the fused statements are adaptively more challenging than the base problems to fit prover performance evolution.

2.1.2 PROOF WRITING

Utilizing the NL-FL statement pair generated by the fusion process, the prover from the last iteration generates n candidate proofs (with $n = 16$ in our implementation), denoted as $\{o_k\}_{k=1}^n$, specifically:

$$o_k = Prover_{i-1}(s^{(NL)}, s^{(FL)}), \forall k \in [1, n]$$

When $i = 0$, it is the base prover model like deepSeek-Prover-V2 (Ren et al., 2025) or Goedel-Prover-V2 (Lin et al., 2025b).

The set of candidate proofs $\{o_k\}_{k=1}^n$ is then passed to the Lean4 verifier to check correctness. Based on the verification results, we compute the empirical pass rate p of this statement for reward assignment as $p = \frac{n_{pass}}{n}$, where n_{pass} is the number of proofs that pass verification.

2.2 ADVERSARIAL REINFORCEMENT LEARNING

This section outlines the training process of **GAR**, where the prover and statement fuser are optimized in competition, deriving adversarial learning and mutual improvement.

2.2.1 STATEMENT FUSER TRAINING

The statement fuser is trained to synthesize problems that are slightly beyond the prover’s current capability, thereby creating an implicit curriculum. To achieve this, we assign rewards to the generated NL statements that encourage the production of more challenging but solvable problems. We adopt a variant of Group Relative Policy Optimization (GRPO) (Shao et al., 2024) for this optimization. Specifically, for each step i , the fuser generates N statements (with $N = 1024$ in our implementation). Formally, generated statement set is: $\{s_{i,j}^{(NL)}\}_{j=1}^N$, and base statement set is represented as:

$\{s_{i,j}^{(base)} = (s_{i,j,1}^{(NL)}, s_{i,j,2}^{(NL)})\}_{j=1}^N$. The statement fuser $Fuser_{i-1}$ is then updated by maximizing the following objective.

$$\mathcal{J}_{GRPO}^{(F)}(\theta) = \mathbb{E} \left[\left\{ \mathbf{s}_{i,j}^{(base)} \right\}_{j=1}^N \sim \mathcal{D}_{stat}; \left\{ \mathbf{s}_{i,j}^{(NL)} \sim \pi_{\theta_{old}}(\mathbf{s}_{i,j}^{(base)}) \right\} \right]$$

$$\frac{1}{N} \sum_{j=1}^N \left(\min \left\{ \delta_{i,j}^{(stat)} A_{i,j}^{(stat)}, \text{clip} \left\{ \delta_{i,j}^{(stat)}, 1 \pm \varepsilon \right\} A_{i,j}^{(stat)} \right\} - \beta \mathbb{D}_{KL}(\pi_{\theta} \parallel \pi_{\theta_{ref}}) \right), \quad (1)$$

$$\delta_{i,j}^{(stat)} = \frac{\pi_{\theta} \left(\mathbf{s}_{i,j}^{(NL)} \mid \mathbf{s}_{i,j}^{(base)} \right)}{\pi_{\theta_{old}} \left(\mathbf{s}_{i,j}^{(NL)} \mid \mathbf{s}_{i,j}^{(base)} \right)}, \quad (2)$$

$$\mathbb{D}_{KL}(\pi_{\theta} \parallel \pi_{\theta_{ref}}) = \frac{\pi_{\theta_{ref}} \left(\mathbf{s}_{i,j}^{(NL)} \mid \mathbf{s}_{i,j}^{(base)} \right)}{\pi_{\theta} \left(\mathbf{s}_{i,j}^{(NL)} \mid \mathbf{s}_{i,j}^{(base)} \right)} - \log \frac{\pi_{\theta_{ref}} \left(\mathbf{s}_{i,j}^{(NL)} \mid \mathbf{s}_{i,j}^{(base)} \right)}{\pi_{\theta} \left(\mathbf{s}_{i,j}^{(NL)} \mid \mathbf{s}_{i,j}^{(base)} \right)} - 1 \quad (3)$$

where π_{θ} denotes the policy model with parameter θ . It represents the statement fuser here. θ_{ref} is the parameter for base statement fuser, θ_{old} is the parameter for the fuser in the previous step, ε and β are hyper-parameters, and $A_{i,j}$ is the advantage, computed from the reward set $\mathbf{r}_i^{(stat)} = \{r_{i,j}^{(stat)}\}_{j=1}^{N_i}$ by:

$$A_{i,j}^{(stat)} = \frac{r_{i,j}^{(stat)} - \text{mean}(\mathbf{r}_i^{(stat)})}{\text{std}(\mathbf{r}_i^{(stat)})}, \quad r_{i,j}^{(stat)} = (1 - p_{i,j}) \cdot (1 - m_{i,j}) \cdot \mathbb{I}\{p_{i,j} \neq 0\} \quad (4)$$

where $p_{i,j}$ is the prover’s empirical pass rate on $s_{i,j}^{(FL)}$, and $m_{i,j}$ is the statement modification rate, which indicates the portion of proofs that the prover tries to modify its statement. The reward is set to 0 if the prover fails to solve the problem, which indicates the statement is too difficult or unsolvable.

We introduce the term $1 - m_{i,j}$ as a soft statement modification penalty to balance the risk between reward hacking and the need to preserve model capability. Because of Long CoT training, current expert provers often acquire strong self-correction capability. While valuable, this ability can lead the model to change formal statements during proof writing. This may lead to serious reward hacking if unconstrained. Conversely, a strict ban on such modifications would suppress self-correction and reduce proofreading accuracy. Our soft penalty discourages excessive statement change without too much harm to the result.

In summary, the training design for the statement fuser rewards the model to lower the prover’s pass rate by composing more challenging statements.

2.2.2 PROVER TRAINING

To achieve adversarial training, the prover is optimized to maximize the pass rate on generated statements. We employ a variant of the GRPO algorithm for such training. Specifically, for each theorem statement $\mathbf{s}_{i,j} = (s_{i,j}^{(NL)}, s_{i,j}^{(FL)})$, we update the prover model $Prover_{i-1}$ by maximizing the following objective function:

$$\mathcal{J}_{GRPO}^{(P)}(\omega) = \mathbb{E} \left[\left\{ o_{i,j,k} \right\}_{k=1}^n \sim \pi_{\omega_{old}}(\mathbf{s}_{i,j}) \right]$$

$$\frac{1}{n} \sum_{k=1}^n \left(\min \left\{ \delta_{i,j,k}^{(pr)} A_{i,j,k}^{(pr)}, \text{clip} \left\{ \delta_{i,j,k}^{(pr)}, 1 \pm \varepsilon \right\} A_{i,j,k}^{(pr)} \right\} - \beta \mathbb{D}_{KL}(\pi_{\omega} \parallel \pi_{\omega_{ref}}) \right), \quad (5)$$

$$\delta_{i,j,k}^{(proof)} = \frac{\pi_{\omega} (o_{i,j,k} \mid \mathbf{s}_{i,j})}{\pi_{\omega_{ref}} (o_{i,j,k} \mid \mathbf{s}_{i,j})} \quad (6)$$

where π_{ω} is the prover as policy model, ω is the parameter of prover, and $A_{i,j,k}$ is the advantage of proof $o_{i,j,k}$, computed from the reward set $\mathbf{r}_{i,j}^{(proof)} = \{r_{i,j,k}^{(proof)}\}_{k=1}^n$ by:

$$A_{i,j,k}^{(proof)} = \frac{r_{i,j,k}^{(proof)} - \text{mean}(\mathbf{r}_{i,j}^{(proof)})}{\text{std}(\mathbf{r}_{i,j}^{(proof)})}, \quad r_{i,j,k} = 1 - 0.5 \cdot m_{i,j,k} \quad (7)$$

where $m_{i,j,k}$ is a binary indicator of whether a statement modification occurred in proof $o_{i,j,k}$. Similar to the training of the statement fuser, modifications are penalized but not strictly prohibited. Furthermore, to ensure the prover is trained on high-quality data, we exclude statements with an empirical pass rate of 0 (unsolvable) or above 0.5 (too easy) following experience in Wang et al. (2025a); Dong & Ma (2025b). It make sure the prover is only trained with hard and medium-level problems. This optimization scheme drives the prover to compete with the statement fuser by continually enhancing its proof-generation capability.

2.3 SUMMARY

The **GAR** is an iterative framework (Wang et al., 2023) that loops the generation stage and the RL stage details above. To provide a clear overview, we present the complete framework in the form of pseudo-code as follows:

Algorithm 1 GAR

Require: \mathcal{D}_{stat} , Statement Fuser (base) as $Fuser$, Prover (base) as $Prover$, Autoformalizer as AF

- 1: $\pi_\omega, \pi_{\omega_{old}}, \pi_{\omega_{ref}} = Prover; \pi_\theta, \pi_{\theta_{old}}, \pi_{\theta_{ref}} = Fuser$
- 2: **for** $i : [1, T]$ **do** ▷ Step i of **GAR**
- 3: $\{\mathbf{s}_{i,j}^{(base)}\}_{j=1}^N \sim \mathcal{D}_{stat}$ ▷ Sample base statements
- 4: $\{\mathbf{s}_{i,j}^{(NL)}\}_{j=1}^N \sim \pi_\theta(\mathbf{s}_{i,j}^{(base)})$ ▷ Fuse NL statements
- 5: $\{\mathbf{s}_{i,j}^{(FL)}\}_{j=1}^N \sim AF(\mathbf{s}_{i,j}^{(NL)})$ ▷ Autoformalize statement
- 6: $\{\mathbf{s}_{i,j} = (\mathbf{s}_{i,j}^{(NL)}, \mathbf{s}_{i,j}^{(FL)})\}_{j=1}^N$
- 7: **for** $j : [1, N]$ **do**
- 8: $\{o_{i,j,k}\}_{k=1}^K \sim \pi_\omega(\mathbf{s}_{i,j})$ ▷ Generate proofs
- 9: Obtain $p_{i,j}, m_{i,j}$ by Lean checking of proofs
- 10: $\pi_\theta = \pi_{\theta_{old}} \leftarrow Optimize(\mathcal{J}_{GRPO}^{(F)}(\theta))$ ▷ Train NL Fuser
- 11: $\pi_\omega = \pi_{\omega_{old}} \leftarrow Optimize(\mathcal{J}_{GRPO}^{(P)}(\omega))$ ▷ Train Prover
- 12: **return** π_θ, π_ω

When trained in multiple iterations, the **GAR** establishes an implicit curriculum by aligning statement difficulty with the prover’s evolving capability. The statement fuser is trained to reduce the pass rate by generating more challenging statements, while the prover is optimized to increase it by producing more valid proofs. Together, their adversarial interaction drives progressive improvement.

3 EXPERIMENTS

We conduct comprehensive experiments on the MiniF2F-Test (Zheng et al., 2021) and ProofNet-Test (Azerbaiyev et al., 2023) benchmarks to assess the performance of the **GAR** framework in formal proof writing. Specifically, we show in Section 3.3 that the models trained with **GAR** achieve better empirical results, confirm in Section 3.4 that adversarial training induces an implicit curriculum by generating progressively harder statements, and report ablation study results in Section 3.5. Due to space limitations, we have included the efficiency study and case study in Appendix D.

3.1 EXPERIMENT SETUP

3.1.1 DATASET AND TASK

We measure the Lean4 reasoning capability of the **GAR** trained model by MiniF2F-Test (Zheng et al., 2021) and ProfNet-Test (Azerbaiyev et al., 2023) benchmarks. They are challenging benchmarks and adopted in nearly all major studies in the field (Xin et al., 2024b; Lin et al., 2024; Wang et al., 2024; Wu et al., 2024a; Polu et al., 2022; Dong & Ma, 2025a; Lin et al., 2025b; Wang et al., 2025c).

The MiniF2F-Test benchmarks comprise 244 Lean4 statements, spanning from high school competition problems to elementary undergraduate-level theorem proofs. Specifically, MiniF2F-Test

Table 1: Main experimental results of **GAR** trained models compared to recent provers models.

Method	Sample budget	MiniF2F-Test	ProofNet-Test
Lean-STaR (Lin et al., 2024)	$64 \times 1 \times 50$	46.31%	-
InternLM-2.5-StepProver (Wu et al., 2024a)	$4 \times 32 \times 600$	50.70%	18.80%
DeepSeek-Prover-V1.5-RL (Xin et al., 2024b)	128	50.00%	18.20%
STP-Lean (Dong & Ma, 2025b)	128	56.15%	19.50%
MA-LoT (Wang et al., 2025c)	32	61.07%	15.47%
Kimina-Prover-Distill-7B (Wang et al., 2025a)	32	63.10%	-
DeepSeek-Prover-V2-7B (Ren et al., 2025)	32	70.49%	22.58%
Goedel-Prover-V2-8B (Lin et al., 2025a)	32	77.87%	-
<i>Our models</i>			
GAR on Deepseek-Prover-V2	32	74.18%	25.81%
GAR on Goedel-Prover-V2	32	80.33%	-

Table 2: Average proof correctness rate for Goedel-Prover-V2-8B (base model) and GAR model trained on base model.

Step idx	Base Model	GAR model
0	29.16%	29.16%
1	16.50%	23.71%
2	11.58%	20.53%
3	7.61%	20.08%
4	7.69%	21.79%

Table 3: Statement modification rate for dropping statement modification penalty and full GAR trained models.

Step idx	w/o Stat. Penalty	Full GAR
0	42.94%	42.94%
1	48.18%	48.72%
2	60.42%	30.50%
3	71.82%	39.65%
4	74.11%	33.63%

comprises problems formalized from the MATH dataset (Hendrycks et al., 2021), high school competitions such as AMC, AIME, and IMO, as well as self-crafted problems. ProofNet-Test consists of 186 theorems formalized from standard undergraduate textbooks on advanced topics such as real and complex analysis, linear algebra, abstract algebra, and topology. In our setting, we train the LLM with **GAR** to generate Lean4 proofs from the NL-FL statement pair. To avoid overloading the model, all the imports and namespaces are manually configured.

3.1.2 BASELINES

To evaluate the effectiveness of **GAR**, we compare it against strong open-source baselines, including Lean-STaR (Lin et al., 2024), InternLM-2.5-StepProver (Wu et al., 2024a), Kimina-Prover-Preview-Distill-7B (Wang et al., 2025a), DeepSeek-Prover-V1.5-RL (Xin et al., 2024b), STP-Lean (Dong & Ma, 2025a), MA-LoT (Wang et al., 2025c), Goedel-Prover-V2 (Lin et al., 2025b), and DeepSeek-Prover-V2 (Ren et al., 2025). For baseline models that require Long CoT reasoning (Kimina-Prover, MA-LoT, DeepSeek-Prover-V2, and Goedel-Prover-V2), we restrict the reasoning length to 16,384 tokens to conserve computational resources.¹

3.2 IMPLEMENTATION DETAILS

In the generation stage, we construct a repository of 793,243 NL statements from Numina-Math (LI et al., 2024) and Lean-Workbook (Ying et al., 2024) datasets. For statement fusion, we employ Qwen3-8B (Team, 2025) as the base model for the statement fuser because of its skip-thinking capability. We apply Kimina-Autoformalizer-7B (Wang et al., 2025a) as the autoformalizer. For proof generation, we use DeepSeek-Prover-V2-7B (Ren et al., 2025) and Goedel-Prover-V2 (Lin et al., 2025b) as base models for provers. We sample 1,024 theorems per step and generate 16 proofs per theorem following (Wang et al., 2025a). We also restrict the sequence length to 16,384 tokens for both models. The GRPO hyperparameters are set with a learning rate of $2E-6$, $\varepsilon = 0.2$, and $\beta = 0.01$. We perform three iterations of **GAR** training on Goedel-Prover-V2 and five iterations on DeepSeek-Prover-V2, which costs around 140 H100 hours for each training. On average, DeepSeek-Prover-V2 costs fewer hours per iteration due to average shorter thinking. In verification, any proof

¹Because of this restriction, our reported results for DeepSeek-Prover-V2 and Goedel-Prover-V2 differ from those in the original papers, where the evaluations used a context length of 40,960.

involving statement modification or relying on the “sorry” or “admit” tactic is counted as incorrect to ensure fairness.

3.3 RESULTS

Table 1 demonstrates the empirical results of applying **GAR** to train base prover models. For Goedel-Prover-V2, MiniF2F-Test pass@32 raises to 80.33%, indicating a **3.16%** relative gain. For DeepSeek-Prover-V2, the MiniF2F-Test score improves to 74.18%, corresponding to a **5.23%** increase. On the more challenging ProofNet-Test benchmark, which targets advanced mathematics topics, **GAR** enhances DeepSeek-Prover-V2’s pass@32 rate from 22.58% to **25.81%**.²

Compared with other baselines, both the base model and the **GAR**-trained models achieve consistent and substantial gains, highlighting that **GAR** can contribute to models that already have outstanding performance. Compared to base models, the improvements suggest that the adversarial method proposed in **GAR** enables provers to tackle increasingly difficult problems by establishing an implicit curriculum. Such a curriculum can progressively adapt problem difficulty to the model’s capability, allowing the prover to explore deeper reasoning strategies and solve more advanced theorems. This behavior becomes more significant as the theorems get harder.

3.4 STATEMENT FUSER STUDY

To show that **GAR** produces increasingly more difficult statements across iterations and thereby establishes an implicit curriculum learning, we examine the difficulty of problems generated by the statement fuser. In each iteration, we randomly sample 50 generated statements and compute their average proof correctness rate using both the base prover and **GAR**-trained prover at that iteration. Details of the metric we used are provided in Appendix B.2. The result of this experiment is reported in Table 2.

From the results, we can observe that the base model exhibits a consistent performance degradation, with accuracy falling from 29.16% at the first iteration’s data to 7.69% by the fifth. In contrast, the **GAR**-trained model shows only a minor initial decline due to the statement matching but stabilizes at around 21% across later iterations. These findings confirm that **GAR** progressively generates harder statements, as evidenced by the base model’s decline. The maintained trained model’s performance demonstrates the adversarial training strength of the prover’s performance over time.

3.5 ABLATION STUDIES

3.5.1 EFFECT OF STATEMENT MODIFICATION PENALTY

We evaluate the impact of the statement modification penalty by monitoring statement modification rates across by the provers during training rounds. Details of this metric are provided in Appendix B.3. We trained a variant of **GAR** on Goedel-Prover-V2-8B without the penalty in both the statement fuser and prover, and compared it with the full **GAR**. Table 3 shows that without the modification penalty, the prover exploits its self-correction ability by simplifying statements. Such behavior worsens as the training progresses; by the fourth step, 74% of statements were modified at least once, which is a clear signal of reward hacking. In contrast, with the penalty enabled, the modification rate remains below 40% throughout. This study confirms the effectiveness of our statement penalty in preventing reward hacking.

Table 4: Pass@32 results on MiniF2F-Test for Goedel-Prover-V2-8B and its GRPO and GAR trained version.

Method	MiniF2F-Test
Base model	77.87%
GRPO trained	77.46%
GAR trained	80.33%

3.5.2 COMPARE TO DIRECT RL TRAINING

This experiment tests whether jointly evolving the prover and statement fuser is more effective than training the prover alone using existing data. We conducted three additional GRPO iterations on

²We do not report Goedel-Prover-V2’s accuracy on ProofNet due to the absence of the reference results in Lin et al. (2025b).

Goedel-Prover-V2-8B using the same sampling budget and formalized NL base dataset as **GAR**. The results are shown in Table 4. The **GAR**-tained model outperforms the variant trained with traditional GRPO. We conclude that such an improvement to **GAR** progressively raises problem difficulty, enabling the prover to handle more complex statements. In contrast, further RL training on static dataset degrades performance, as the base model is already heavily RL trained. These findings further support the generality of **GAR**, demonstrating its ability to enhance model’s performance while standard RL no longer provides benefits.

4 RELATED WORK

4.1 LLM FOR FORMAL THEOREM PROVING

Recently, applying LLMs to support formal theorem proving has become a prominent research direction. Training approaches can be broadly divided into two categories: model trained solely with Supervised Fine-Tuning (SFT) and those combining SFT with Reinforcement Learning (RL). Early SFT-based provers include Expert Iteration (Polu & Sutskever, 2020), Re-Prover (Yang et al., 2024c), TheoremLlama (Wang et al., 2024), DeepSeek-Prover-V1 (Xin et al., 2024a), InternLM-2.5-StepProver (Wu et al., 2024a), MA-LoT (Wang et al., 2025c), and Goedel-Prover-V1 (Lin et al., 2025a). These models typically require multiple rounds of large-scale annotation with existing formal solvers, which demands a significant amount of computational resources and limits exploration. To further advance the prover, researchers began incorporating advanced RL techniques. For instance, DeepSeek-Prover-V1.5 (Xin et al., 2024b) employs DPO. After the “ZERO” RL technique that enables Long CoT thinking developed by Guo et al. (2025), provers like Kimina-Prover (Wang et al., 2025a), DeepSeek-Prover-V2 (Ren et al., 2025), and Goedel-Prover-V2 (Lin et al., 2025b) enable models to produce better formal proof after thinking. However, existing RL approaches still rely on a fixed collection of statements, preventing statements from adapting to the prover’s evolving skill. In contrast, **GAR** jointly trains a statement fuser and a prover, ensuring that generated theorems remain suitably challenging as the model improves.

4.2 RL METHODS FOR LLM

RL has been central to the development of reasoning LLMs (Xiong et al., 2025). Early efforts such as DeepSeek-Math (Shao et al., 2024) and Qwen-2.5-Math (Yang et al., 2024a) applied reward modeling with GRPO algorithm to enhance exploration. The release of OpenAI-O1 (OpenAI, 2024) and DeepSeek-R1 (Guo et al., 2025) demonstrated that verifier-based rewards in RL can enable complex reasoning strategies, including backward search and self-correction. Such techniques have been widely adopted in systems like Qwen3 (Team, 2025) and Kimi-K2 (Team et al., 2025). Nevertheless, nearly all prior work uses the verification signal only to optimize the problem solver, leaving the problem composer untrained. On the other hand, **GAR** simultaneously improves both statement fuser and the prover, creating an implicit curriculum that adapts task difficulty and allows the model to acquire more complex reasoning skills step-by-step.

5 CONCLUSION

This paper presents **GAR**, a comprehensive training framework for formal theorem proving. **GAR** aims to resolve the inefficient and suboptimal performance caused by traditional expert iteration and online RL in prover training. **GAR** achieves more efficient training by using statement fusion to formulate statements in the RL process to avoid the prover from annotating problems beyond its capability range. After **GAR** training, provers can solve more advanced problems through the implicit curriculum learning that enables the model to explore further. Furthermore, **GAR** can improve the performance of models that have been through heavy RL training by letting the model explore more difficult statements during the training process. Experiments of applying **GAR** to DeepSeek-Prover-V2 and Goedel-Prover-V2 yield an average relative improvement of **4.20%** on the MiniF2F-Test dataset and improve the DeepSeek-Prover-V2’s performance on ProofNet-Test from 22.58% to **25.81%**. Beyond theorem proving, **GAR** offers a general RL paradigm of co-evolution of the problem generation and solving under a verifiable environment. It provides a foundation for adversarial co-training in other reasoning-intensive domains.

ETHICS STATEMENT

After carefully reviewing the ethical regulations of the conference, to the best of our knowledge, this work does not present any foreseeable ethical concerns. This research focuses exclusively on formal theorem proving using publicly available mathematical datasets, without involving human subjects, private data, or sensitive content. As far as we are concerned, no negative societal or ethical impacts are anticipated for the contribution of this work. We only use LLMs to polish the writing style and fix grammatical errors in the paper.

REPRODUCIBILITY STATEMENT

We have made efforts to ensure that our work is reproducible. The detailed description of the **GAR** framework, including pseudocode, data source, reward definitions, and training objectives, is presented in Section 2. The experimental setup, including benchmarks, baselines, base model choice, and hyperparameters, is presented in Section 3.1 and Section 3.2. We also plan to further ensure the reproducibility by open-sourcing the code in the near future.

REFERENCES

- 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.
- Projet Coq. The coq proof assistant-reference manual. *INRIA Rocquencourt and ENS Lyon, version, 5*, 1996.
- Leonardo De Moura, Soonho Kong, Jeremy Avigad, Floris Van Doorn, and Jakob von Raumer. The lean theorem prover (system description). In *Automated Deduction-CADE-25: 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings 25*, pp. 378–388. Springer, 2015.
- Kefan Dong and Tengyu Ma. Beyond limited data: Self-play llm theorem provers with iterative conjecturing and proving. *arXiv preprint arXiv:2502.00212*, 2025a.
- Kefan Dong and Tengyu Ma. Stp: Self-play llm theorem provers with iterative conjecturing and proving. *arXiv preprint arXiv:2502.00212*, 2025b.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- John Harrison. Hol light: An overview. In *International Conference on Theorem Proving in Higher Order Logics*, pp. 60–66. Springer, 2009.
- 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. *arXiv preprint arXiv:2103.03874*, 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. *arXiv preprint arXiv:2210.12283*, 2022.
- Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang, Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann Fleureau, Guillaume Lample, and Stanislas Polu. NuminaMath. [<https://huggingface.co/AI-MO/NuminaMath-CoT>] (https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf), 2024.
- Yuxin Li, Minghao Liu, Ruida Wang, Wenzhao Ji, Zhitao He, Rui Pan, Junming Huang, Tong Zhang, and Yi R Fung. Lean4physics: Comprehensive reasoning framework for college-level physics in lean4. *arXiv preprint arXiv:2510.26094*, 2025.

- Haohan Lin, Zhiqing Sun, Yiming Yang, and Sean Welleck. Lean-star: Learning to interleave thinking and proving. *arXiv preprint arXiv:2407.10040*, 2024.
- 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, 2025a. URL <https://arxiv.org/abs/2502.07640>.
- Yong Lin, Shange Tang, Bohan Lyu, Ziran Yang, Jui-Hui Chung, Haoyu Zhao, Lai Jiang, Yihan Geng, Jiawei Ge, Jingruo Sun, et al. Goedel-prover-v2: Scaling formal theorem proving with scaffolded data synthesis and self-correction. *arXiv preprint arXiv:2508.03613*, 2025b.
- 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 28*, pp. 625–635. Springer, 2021.
- Allen Newell and Herbert Simon. The logic theory machine—a complex information processing system. *IRE Transactions on information theory*, 2(3):61–79, 1956.
- OpenAI. Learning to reason with llms. <https://openai.com/index/learning-to-reason-with-llms/>, September 13 2024. Accessed: 2024-11-24.
- Lawrence C Paulson. *Isabelle: A generic theorem prover*. Springer, 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. *arXiv preprint arXiv:2202.01344*, 2022.
- ZZ Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanxia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, et al. Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition. *arXiv preprint arXiv:2504.21801*, 2025.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025.
- Qwen Team. Qwen3, April 2025. URL <https://qwenlm.github.io/blog/qwen3/>.
- Haiming Wang, Mert Unsal, Xiaohan Lin, Mantas Baksys, Junqi Liu, Marco Dos Santos, Flood Sung, Marina Vinyes, Zhenzhe Ying, Zekai Zhu, et al. Kimina-prover preview: Towards large formal reasoning models with reinforcement learning. *arXiv preprint arXiv:2504.11354*, 2025a.
- Ruida Wang, Wangchunshu Zhou, and Mrinmaya Sachan. Let’s synthesize step by step: Iterative dataset synthesis with large language models by extrapolating errors from small models. *arXiv preprint arXiv:2310.13671*, 2023.
- Ruida Wang, Jipeng Zhang, Yizhen Jia, Rui Pan, Shizhe Diao, Renjie Pi, and Tong Zhang. Theorem-llama: Transforming general-purpose llms into lean4 experts. *arXiv preprint arXiv:2407.03203*, 2024.
- Ruida Wang, Yuxin Li, Yi R Fung, and Tong Zhang. Let’s reason formally: Natural-formal hybrid reasoning enhances llm’s math capability. *arXiv preprint arXiv:2505.23703*, 2025b.
- Ruida Wang, Rui Pan, Yuxin Li, Jipeng Zhang, Yizhen Jia, Shizhe Diao, Renjie Pi, Junjie Hu, and Tong Zhang. Ma-lot: Model-collaboration lean-based long chain-of-thought reasoning enhances formal theorem proving. *arXiv preprint arXiv:2503.03205*, 2025c.

- Yuxiang Wei, Zhe Wang, Jiawei Liu, Yifeng Ding, and Lingming Zhang. Magicoder: Empowering code generation with oss-instruct. *arXiv preprint arXiv:2312.02120*, 2023.
- 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. *arXiv preprint arXiv:2410.15700*, 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, Daya Guo, Zhihong Shao, Zhizhou Ren, Qihao Zhu, Bo Liu, Chong Ruan, Wenda Li, and Xiaodan Liang. Deepseek-prover: Advancing theorem proving in llms through large-scale synthetic data. *arXiv preprint arXiv:2405.14333*, 2024a.
- Huajian Xin, ZZ Ren, Junxiao Song, Zhihong Shao, Wanbiao Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, et al. Deepseek-prover-v1. 5: Harnessing proof assistant feedback for reinforcement learning and monte-carlo tree search. *arXiv preprint arXiv:2408.08152*, 2024b.
- Ran Xin, Chengguang 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. *arXiv preprint arXiv:2502.03438*, 2025.
- Wei Xiong, Jiarui Yao, Yuhui Xu, Bo Pang, Lei Wang, Doyen Sahoo, Junnan Li, Nan Jiang, Tong Zhang, Caiming Xiong, et al. A minimalist approach to llm reasoning: from rejection sampling to reinforce. *arXiv preprint arXiv:2504.11343*, 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*, 2024a.
- Kaiyu Yang, Gabriel Poesia, Jingxuan He, Wenda Li, Kristin Lauter, Swarat Chaudhuri, and Dawn Song. Formal mathematical reasoning: A new frontier in ai. *arXiv preprint arXiv:2412.16075*, 2024b.
- 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. *Advances in Neural Information Processing Systems*, 36, 2024c.
- Jiarui Yao, Ruida Wang, and Tong Zhang. Fans—formal answer selection for natural language math reasoning using lean4. *arXiv preprint arXiv:2503.03238*, 2025.
- 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. *arXiv preprint arXiv:2406.03847*, 2024.
- Hanning Zhang, Ruida Wang, Rui Pan, Wenyuan Wang, Bingxu Meng, and Tong Zhang. Physprover: Advancing automatic theorem proving for physics. *arXiv preprint arXiv:2601.15737*, 2026.
- Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. Minif2f: a cross-system benchmark for formal olympiad-level mathematics. *arXiv preprint arXiv:2109.00110*, 2021.

A ACKNOWLEDGMENT

This material is based upon work supported partially by NSF under Grant No. 2416897, Grant No. 2505932, and by ORN under Grant No. N000142512318. This research used both Delta (NSF award OAC 2005572) and DeltaAI (NSF award OAC 2320345) advanced computing systems, and computing resources provided by Illinois Computes and NAIRR Pilot NAIRR250157.

B DETAILS OF THE METRICS

For completeness, this section provides the metric definitions that were omitted due to space limitations in Section 3.

B.1 PASS RATE

Pass@ x is a widely used metric for evaluating formal theorem provers (Polu et al., 2022; Jiang et al., 2022; Wang et al., 2024; Xin et al., 2024a; Dong & Ma, 2025b; Wang et al., 2025b; Lin et al., 2025b; Wu et al., 2024a; Wang et al., 2025a; Li et al., 2025; Yao et al., 2025; Zhang et al., 2026). For each theorem statement, the model generates x candidate proofs, which are further checked by the Lean4 verifier. If at least one of the x candidates is correct, the theorem is counted as a proved theorem. The pass@ x of the LLM is the fraction of theorems in the dataset with at least one correct proof generated by the prover.

B.2 AVERAGE PROOF CORRECTNESS RATE

This metric measures the relative difficulty of a set of statements for a given prover, which is used in Section 3.4). Let \mathcal{D} be a dataset of size m . For each theorem $t_i \in \mathcal{D}$, we sample 16 proofs from the prover and record the number of proofs that pass the verification, which is p_i . The average proof correctness rate is calculated by

$$\frac{\sum_{i=1}^m p_i}{16 \cdot m}$$

A high value indicates the dataset is easier for the given prover.

Note that in this experiment, to demonstrate the significance of **GAR** better, we train two extra rounds of Goedel-Prover-V2-8B model.

B.3 STATEMENT MODIFICATION RATE

Since current advanced provers have Long Chain-of-thought (CoT) capability, which enables self-reflection, backtracking, and self-correction. In the Long CoT process, a prover may alter the original statement into a simplified variant in its reasoning trace and ultimately produce a proof for the modified version. To quantify this behavior, we define statement modification rate as the portion of theorems in a dataset that have at least one statement modification across 16 generated proofs. This metric helps assess how often the prover attempts to modify the problem and whether such a behavior may be too significant to cause severe reward hacking.

C DIRECT COMPARISON WITH RELATED WORKS

This section aims to provide a more direct comparison between **GAR** and other works that also try to build a dynamic dataset during training, namely Goedel-Prover-V2 (Lin et al., 2025b) and STP (Dong & Ma, 2025b).

Compared to Goedel-Prover-V2: The statement generation in Goedel-V2 relies on a frozen large model to synthesize data for SFT. Crucially, during the RL phase, their statement set remains static. This lack of dynamic updates of statements based on the prover’s evolving capabilities may lead to suboptimal performance and efficiency as the prover outpaces the fixed problem set. In contrast, **GAR** continuously updates the statement fuser via adversarial RL, ensuring the generated statements

Table 5: pass@32 results on PutnamBench

Model Type	PutnamBench
DeepSeek-Prover-V2-7B (Ren et al., 2025)	22/660
GAR DeepSeek-Prover	24/660

Table 6: Statement Modification Rate for MiniF2F-Test under pass@32

Model Type	Base Model	GAR trained
DeepSeek-Prover-V2 (Ren et al., 2025)	6.96%	13.11%
Goedel-Prover-V2 (Lin et al., 2025b)	24.18%	27.05%

remain progressively challenging and aligned with the prover’s current skill level during the RL phase, leading to better empirical performance.

Compared to STP: While STP trains a conjecture model based on the prover’s feedback, it relies on the expert iteration framework based on offline SFT. This process is inherently inefficient because it requires generating a massive volume of training data in a single iteration to achieve effective SFT training. According to Dong & Ma (2025b), they generate 75,000 conjectures per iteration, which is larger than the entire statement set for our generation. Conversely, GAR operates within an online RL cycle. This allows the fuser and prover to continuously update more efficiently, achieving superior performance without the computational burden of the large-scale per-iteration data required by STP.

D ADDITIONAL EXPERIMENTS

This appendix section provides additional experiments that are omitted in the main paper due to space limitations.

D.1 RESULTS ON PUTNAMBENCH

Given that Lean4 formal reasoning is a rapidly progressing field, the MiniF2F-Test serves as a benchmark with high accuracy in current models. To further validate that the **GAR** training can also make the model perform better on more advanced benchmarks, we evaluate both DeepSeek-Prover-V2-7B and GAR-trained DeepSeek-Prover on PutnamBench under pass@32. The results are demonstrated in Table 5.

We can see that the GAR-trained model solves four additional problems compared to the base model on this challenging benchmark. This consistent improvement across MiniF2F, PutnamBench, and ProofNet demonstrates the robustness of the implicit curriculum established by **GAR**.

D.2 EFFICIENCY STUDY

We analyze the computational efficiency of the **GAR** framework, as training from scratch demands substantial resources, to evaluate this, we compare our approach with Kimina-Prover (Wang et al., 2025a), selected for its transparent reporting of training details. The number of roll-out theorems and proofs per iteration in Wang et al. (2025a) matches our configuration. From Figure 4 in their work, we observe that Kimina-Prover achieves approximately 2% of performance gains after 25 training iterations, where **GAR** trained Goedel-Prover-V2-8B reaches more improvements in merely three iterations.

Notably, despite Kimina-Prover starting from an SFT base model, while **GAR** starts from a heavily RL-trained base model, this comparison strengthens our efficiency claim. Because it is well established in prior work (Guo et al., 2025) that models already heavily optimized via RL face diminishing returns and are inherently harder to improve than SFT baselines. The fact that **GAR** achieves relatively significant gains on top of RL-optimized base models in only 3-5 iterations, while Kimina-Prover gains approximately 2% from an SFT starting point over 25 iterations, demonstrates the high sample efficiency of our adversarial training paradigm. Given resource constraints that

prevent exhaustive experimentation, we leave the discussion of the scalability of **GAR** for future work.

D.3 STATEMENT MODIFICATION RATE STUDY

To dive deeper into the behavior of statement modification, we provide detailed statement modification rates of both the base model and the **GAR** trained model in MiniF2F-Test under pass@32. The results are shown in Table 6

From the comparison of base models, we can see that the Goedel’s increased statement modification rate happens together with its performance enhancement. We attribute this to the stronger self-correction capability it obtains. Similarly, the soft penalty in GAR ensures that the model is penalized if it simplifies the problem to a triviality. Furthermore, we can observe that when the modification rate is low, the GAR training introduces a higher modification rate, as is the case for DeepSeek-Prover. However, if the modification rate is high in the base model, the soft penalty will effectively control it within a reasonable range. These findings prove the improvement of the **GAR**-trained model from another point of view

D.4 ADDITIONAL ABLATION STUDIES

D.4.1 FROZEN FUSER STUDY

To further analyze our generative adversarial training, we conduct this experiment that applies the GAR training only on the prover and keeps the statement fuser untrained. We run three iterations of Frozen fuser GAR training on Goedel-Prover-V2-8B for three iterations. The MiniF2F-Test pass@32 results are demonstrated in Table 7.

From the results, we can see that frozen fuser fails to obtain any performance improvements compared to the base model. This confirms that a static generator is unable to extend the prover’s capabilities beyond its initial limits. This proves the necessity of co-evolution for both the problem composer and the prover.

D.4.2 SINGLE PROBLEM ENHANCEMENT FUISER

In order to further study the effectiveness of our idea of problem fusing, which enhances the problem difficulty by combining two problems into a single one. We conduct the experiment of replacing the statement fuser with MagicCoder (Wei et al., 2023) that enhances the problem difficulty by restating the provided single problem to make it more advanced. The results are presented in Table 7.

The result indicates that replacing the fusion process with the MagicCoder style problem-evolving strategy results in a 2.46% performance drop compared to the base model. We conclude this degradation to the limited capability of Qwen3-8B to create more advanced problems with only one reference. It may not provide sufficient difficulty gradient for the prover, leading to training failure. In contrast, our fusion approach creates a more effective implicit curriculum by combining distinct concepts from multiple problems, thereby validating the design of the Statement Fuser.

D.5 CASE STUDY

D.5.1 CASE 1:

In this section, we demonstrate an example of how the **GAR** trained model successfully solves the problem while the base models fail. Compared to the base prover, the **GAR** trained model relies less on brute simplification and more on early structural framing, producing shorter, steadier proof traces that remain faithful to the statement and are more likely to pass the verification. If we dig deeper into the Long CoT of two models, we can find that **GAR**-trained prover writes a shorter and more focused long CoT. While the base model’s CoT is longer but less

Table 7: Additional Ablation Studies

Method	MiniF2F-Test
Base model	77.87%
Frozen Fuser	77.87%
MagicCoder Fuser	75.41%
GAR trained	80.33%

directed, it cycles among general identities and simplifications without stabilizing on a closing invariant. The result that GAR's trace is both clearer and more verifiable, which leads to the successful proof.

GAR-trained model generated result:

<|begin_of_sentence|><|User|>Complete the following Lean 4 code:

```

`lean4
import Mathlib
import Aesop

set_option maxHeartbeats 0

open BigOperators Real Nat Topology Rat

/-- Consider a prime $p$ of the form $4t+1$. Show that $a$ is a
  ↪ primitive root modulo $p$ iff $-a$ is a primitive root modulo $p$. -/
theorem exercise_4_4 {p t: ℕ} (hp0 : p.Prime) (hp1 : p = 4*t + 1)
  (a : ZMod p) :
  IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
  ...

```

Before producing the Lean 4 code to formally prove the given theorem,
 ↪ provide a detailed proof plan outlining the main proof steps and
 ↪ strategies.

The plan should highlight key ideas, intermediate lemmas, and proof
 ↪ structures that will guide the construction of the final formal
 ↪ proof.

<|Assistant|>### Detailed Proof

Problem Analysis:

We are given a prime `p` of the form `p = 4 * t + 1` for some natural
 ↪ number `t`, and an element `a` in the finite field `ZMod p` (where `p`
 ↪ is the prime number). We need to prove that `a` is a primitive root
 ↪ modulo `p` if and only if `-a` is a primitive root modulo `p`.

Recall that `IsPrimitiveRoot a p` means that `a` has multiplicative order
 ↪ `p - 1 = 4 * t` in the multiplicative group `(ZMod p)^x`.

Key Observations:

1. The multiplicative group `(ZMod p)^x` has order `p - 1 = 4 * t`.
2. The negation map `x ↦ -x` is an automorphism of `(ZMod p)^x` because
 ↪ `p` is an odd prime (`p = 4 * t + 1 > 2`, and thus `p` is odd). This
 ↪ means that `-x` has the same order as `x`.
3. The multiplicative order of `a` and `-a` are the same if and only if `a`
 ↪ = -a, but this is not the case unless `p` divides `2 * a`, which is
 ↪ false here because `a` is arbitrary and `p` is prime. However, the
 ↪ actual condition is that `IsPrimitiveRoot a p` if and only if
 ↪ `IsPrimitiveRoot (-a) p` because the order of `-a` is the same as the
 ↪ order of `a` when `p` is of the form `4 * t + 1`.

But actually, we need a better understanding of the multiplicative group
 ↪ `(ZMod p)^x`.

For any odd prime `p`, we have the following properties:

1. The map `x ↦ -x` is an isomorphism on `(ZMod p)^x` because `p` is odd.
 - The inverse is `x ↦ -x`, so the map is a permutation of the
 ↪ multiplicative group.
 2. For any `x` in `(ZMod p)^x`, `x ^ 2` is a square (i.e., `x ^ 2` is a
 ↪ square in `(ZMod p)^x`).
- To see this, observe that `(a * b)^2 = a^2 * b^2` and `-1` is a
 ↪ quadratic residue modulo `p`.

But for our problem, we don't need these general observations. Instead,
 \hookrightarrow we can use the fact that the order of $-a$ is the same as the order
 \hookrightarrow of a .

****Why are the orders the same?***

Since p is an odd prime ($p = 4 * t + 1 > 2$), the multiplicative group
 $\hookrightarrow (ZMod\ p)^x$ is cyclic of order $p - 1 = 4 * t$. Let g be a generator
 \hookrightarrow of $(ZMod\ p)^x$.

Then:

1. If a is a primitive root (`IsPrimitiveRoot a p` is true), then a
 \hookrightarrow has order $4 * t = p - 1$, so $-a$ must also have order $4 * t$
 \hookrightarrow because $-a$ is a unit in $ZMod\ p$ (p is prime, so $-a \neq 0$).
 - However, this is not correct unless -1 has order 2 because $-a =$
 $\hookrightarrow (-1) * a$, and the order of a is $4 * t$, and the order of -1
 \hookrightarrow is 2 if $t \neq 0$ (because $(-1)^2 = 1$). Thus, the order of $-a$
 \hookrightarrow is $\text{lcm}(2, \text{order of } a)$.

Alternatively, a better approach is to use the following fact:

- The map $x \mapsto -x$ is an automorphism of the group $(ZMod\ p)^x$ because
 $\hookrightarrow p$ is odd ($(-1)^2 = 1$ and multiplication is invertible).

****Conclusion:***

Since the map $x \mapsto -x$ is an automorphism of $(ZMod\ p)^x$ (because p is
 \hookrightarrow an odd prime), a has order $4 * t$ if and only if $-a$ has order 4
 $\hookrightarrow * t$. This is because the automorphism $x \mapsto -x$ preserves the order
 \hookrightarrow of elements.

Thus, `IsPrimitiveRoot a p` (`order a = 4 * t`) is equivalent to
 \hookrightarrow `IsPrimitiveRoot (-a) p` (`order (-a) = 4 * t`).

Therefore, the result follows.

Step-by-Step Abstract Plan

1. ****Understand the Multiplicative Group**:**
 - The multiplicative group $(ZMod\ p)^x$ is cyclic of order $p - 1 = 4 * t$
 $\hookrightarrow t$.
 - Any element has a multiplicative order dividing $4 * t$.
2. ****Understand the Automorphism $x \mapsto -x$ **:**
 - The map $x \mapsto -x$ is an automorphism of $(ZMod\ p)^x$ because p is an
 \hookrightarrow odd prime ($(-1)^2 = 1$).
3. ****Order Preservation**:**
 - The order of $-a$ is the same as the order of a , because $x \mapsto -x$
 \hookrightarrow is an automorphism.
4. ****Conclusion**:**
 - a is a primitive root (`order a = p - 1`) iff $-a$ is a primitive
 \hookrightarrow root (`order (-a) = p - 1`).

Lean 4 `have` Statements

```

'''lean4
theorem exercise_4_4 {p t : N} (hp0 : p.Prime) (hp1 : p = 4*t + 1)
  (a : ZMod p) :
  IsPrimitiveRoot a p  $\leftrightarrow$  IsPrimitiveRoot (-a) p := by
  have h_main : IsPrimitiveRoot a p  $\leftrightarrow$  IsPrimitiveRoot (-a) p := by
    sorry
  sorry
'''

```

Complete Lean 4 Proof

```

`lean4
theorem exercise_4_4 {p t: ℕ} (hp0 : p.Prime) (hp1 : p = 4*t + 1)
(a : ZMod p) :
IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
have h_main : IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
  haveI : Fact p.Prime := ⟨hp0⟩
  have h1 : p ≠ 0 := hp0.ne_zero
  have h2 : p ≠ 1 := hp0.ne_one
  have h3 : (-a : ZMod p) = -a := rfl
  constructor
  · -- Prove the forward direction: if a is a primitive root, then -a
    ↪ is a primitive root
    intro h
    -- Use the property that the order of -a is the same as the order
    ↪ of a
    have h4 := h.isUnit
    have h5 := h.pow_eq_one
    have h6 := h.pow_ne_one_of_pos_of_lt (by decide : (0 : ℕ) < 1) (by
      rw [← Nat.cast_one]
      simp_all [ZMod.val_one, Nat.cast_add, Nat.cast_mul,
        ↪ Nat.cast_ofNat]
      <|> omega)
    simp_all [IsPrimitiveRoot, ZMod.nat_cast_self, Nat.cast_add,
      ↪ Nat.cast_mul, Nat.cast_ofNat,
        ZMod.nat_cast_self]
    <|>
    (try omega) <|>
    (try simp_all [ZMod.nat_cast_self, Nat.cast_add, Nat.cast_mul,
      ↪ Nat.cast_ofNat,
        ZMod.nat_cast_self]) <|>
    (try
      {
        aesop
      }) <|>
    (try
      {
        simp_all [Units.ext_iff, pow_succ]
        <|>
        ring_nf at * <|>
        omega
      }) <|>
    (try
      {
        norm_num at *
        <|>
        aesop
      }) <|>
    (try
      {
        aesop
      }) <|>
    (try
      {
        simp_all [Units.ext_iff, pow_succ]
        <|>
        ring_nf at * <|>
        omega
      })
    <|>
    aesop
  · -- Prove the reverse direction: if -a is a primitive root, then a
    ↪ is a primitive root
    intro h

```

```

-- Use the property that the order of -a is the same as the order
↪ of a
have h4 := h.isUnit
have h5 := h.pow_eq_one
have h6 := h.pow_ne_one_of_pos_of_lt (by decide : (0 : ℕ) < 1) (by
  rw [+ Nat.cast_one]
  simp_all [ZMod.val_one, Nat.cast_add, Nat.cast_mul,
    ↪ Nat.cast_ofNat]
  <|> omega)
simp_all [IsPrimitiveRoot, ZMod.nat_cast_self, Nat.cast_add,
  ↪ Nat.cast_mul, Nat.cast_ofNat,
  ZMod.nat_cast_self]
<|>
(try omega) <|>
(try simp_all [ZMod.nat_cast_self, Nat.cast_add, Nat.cast_mul,
  ↪ Nat.cast_ofNat,
  ZMod.nat_cast_self]) <|>
(try
  {
    aesop
  }) <|>
(try
  {
    simp_all [Units.ext_iff, pow_succ]
    <|>
    ring_nf at * <|>
    omega
  }) <|>
(try
  {
    norm_num at *
    <|>
    aesop
  }) <|>
(try
  {
    aesop
  }) <|>
(try
  {
    simp_all [Units.ext_iff, pow_succ]
    <|>
    ring_nf at * <|>
    omega
  })
<|>
aesop
exact h_main
...

```

Base model generated result:

<|begin_of_sentence|><|User|>Complete the following Lean 4 code:

```

...lean4
import Mathlib
import Aesop

set_option maxHeartbeats 0

open BigOperators Real Nat Topology Rat

```

```

/-- Consider a prime $p$ of the form $4t+1$. Show that $a$ is a
↪ primitive root modulo $p$ iff $-a$ is a primitive root modulo $p$. -/
theorem exercise_4_4 {p t: ℕ} (hp0 : p.Prime) (hp1 : p = 4*t + 1)
  (a : ℤMod p) :
  IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
  ...

```

Before producing the Lean 4 code to formally prove the given theorem,
 ↪ provide a detailed proof plan outlining the main proof steps and
 ↪ strategies.

The plan should highlight key ideas, intermediate lemmas, and proof
 ↪ structures that will guide the construction of the final formal
 ↪ proof.

<|Assistant|>### Detailed Proof and Analysis

First, let's recall the definitions and properties we will use:

1. **Primitive Root**: An element a is a primitive root modulo p if
 ↪ the multiplicative order of a in $(\mathbb{Z}/p\mathbb{Z})^*$ is $p - 1$.
2. **Order of a** : The order of a modulo p , denoted $\text{orderOf } a$, is
 ↪ the smallest positive integer k such that $a^k \equiv 1 \pmod{p}$.
3. **Multiplicative Order**: The multiplicative order of a modulo p is
 ↪ the smallest positive integer k such that $a^k \equiv 1 \pmod{p}$.

Given p is prime and $p = 4t + 1$, we need to prove that a is a
 ↪ primitive root modulo p if and only if $-a$ is a primitive root
 ↪ modulo p .

Key Observations:

1. The multiplicative group $(\mathbb{Z}/p\mathbb{Z})^*$ has order $p - 1 = 4t$.
2. The quadratic residues modulo p are $\{a \mid \exists b, b^2 \equiv a \pmod{p}\}$. The
 ↪ number of quadratic residues is $(p - 1)/2 = 2t$ (since $p \equiv 1 \pmod{4}$).
 ↪ 4).
3. The multiplicative order of any $a \in (\mathbb{Z}/p\mathbb{Z})^*$ must divide $p - 1 =$
 ↪ $4t$.
4. The multiplicative order of $-a$ is the same as the multiplicative
 ↪ order of a because $(-a)^2 = a^2 \equiv a^2 \pmod{p}$, and the order is
 ↪ determined by the exponentiation.

Proof Sketch:

1. Assume a is a primitive root modulo p . This means that a has
 ↪ multiplicative order $p - 1 = 4t$ in $(\mathbb{Z}/p\mathbb{Z})^*$.
 - The order of $-a$ is the same as the order of a , which is $4t = p$
 ↪ $- 1$, because:
 - $(-a)^{4t} = ((-a)^2)^{2t} = (a^2)^{2t} \equiv 1 \pmod{p}$ because a^2
 ↪ has order $2t$ (since a is a primitive root and the exponent
 ↪ $4t = 2 \cdot (2t)$, but this is not directly relevant).
 - More precisely, $-a$ and a have the same multiplicative order
 ↪ because $(-a)^{4t} = ((-a)^2)^{2t} = (a^2)^{2t} \equiv 1 \pmod{p}$ (if
 ↪ a^2 has order $2t$).
 - However, a simpler approach is to observe that a^2 has order $2t$
 ↪ and thus $(-a)^2 = a^2$, so $(-a)^{2t} = (a^2)^t \equiv 1 \pmod{p}$ is not
 ↪ guaranteed. Instead, we can use that $(-a)^{4t} = ((-a)^2)^{2t}$
 ↪ $= (a^2)^{2t} \equiv 1 \pmod{p}$, but this does not directly help unless we
 ↪ know a^2 has order $2t$.
 - A better approach is to note that $(-a)^k \equiv 1 \pmod{p}$ if and only if
 ↪ $a^k \equiv 1 \pmod{p}$ because:
 - $(-a)^k \equiv 1 \pmod{p}$ is equivalent to $(-1)^k * a^k \equiv 1 \pmod{p}$, so
 ↪ $a^k \equiv (-1)^k \pmod{p}$.
 - Since a is a primitive root, $a^k \equiv 1 \pmod{p}$ if and only if k
 ↪ is a multiple of $4t$, because a has order $4t$.

- Thus, $a^k \equiv (-1)^k \pmod{p}$ would imply $(-1)^k \equiv 1 \pmod{p}$, i.e.,
 - ↪ $(-1)^{4t} \equiv 1 \pmod{p}$, which is true for $p = 4t + 1$
 - ↪ $(-1)^{4t+1} \equiv -1 \equiv 1 \pmod{p}$? Wait no: $p = 4t + 1$ and $p \equiv 1 \pmod{4}$, so $(-1)^{4t+1} = (-1)^{4t} * (-1) = 1 * (-1) = -1 \equiv p - 1 \pmod{p}$, which is not $1 \pmod{p}$. But p is prime and $p \equiv 1 \pmod{4}$, so $-1 \equiv p - 1 \equiv 4t \equiv 0 \pmod{4}$, which is not helpful. This is a mistake in the reasoning.
- The correct approach is to note that $(-a)^{4t} = ((-a)^2)^{2t} = (a^2)^{2t} \equiv 1 \pmod{p}$ because a^2 has order $2t$. But a is a primitive root, so $a^{2t} \equiv 1 \pmod{p}$ ($\text{orderOf } a = 4t \geq 2t$).
 - ↪ This is not sufficient.
- Alternatively, observe that:
 - The order of $-a$ divides $4t$ because the order of a divides $4t$.
 - The order of $-a$ does not divide $2t$ unless $a^2 \equiv 1 \pmod{p}$ (since $(-a)^2 = a^2 \equiv 1 \pmod{p}$ would imply $\text{orderOf } (-a) \leq 2$).
 - ↪ But a is a primitive root, so $a^2 \equiv 1 \pmod{p}$ would imply $2 \leq \text{orderOf } a = 4t$, i.e., $2 \leq 4t$, which is true since $t \geq 1$.
 - ↪ But $a^2 \equiv 1 \pmod{p}$ is not guaranteed unless $a^2 \equiv 1 \pmod{p}$, which is not necessarily true.
- A better way is to use the fact that $-a$ has the same order as a because $p \equiv 1 \pmod{4}$, and $(-a)^2 = a^2$.

However, a simpler and correct approach is to use the following:

- The multiplicative order of $-a$ modulo p is the same as the multiplicative order of a modulo p , because $(-a)^2 = a^2$ and a^2 has order $2t$ modulo p .

But we can use the following:

- If a is a primitive root modulo p , then for any k , $a^k \equiv 1 \pmod{p}$ if and only if k is a multiple of $4t$.
- Then $(-a)^k \equiv 1 \pmod{p}$ is equivalent to $a^k \equiv 1 \pmod{p}$, because $(-a)^k = (-1)^k * a^k$ and p is odd ($p \equiv 1 \pmod{4}$). So $(-a)^k \equiv 1 \pmod{p}$ is equivalent to $a^k \equiv 1 \pmod{p}$ if p is odd.
- Hence, $\text{orderOf } (-a) = \text{orderOf } a = 4t = p - 1$.

Therefore, $-a$ is also a primitive root modulo p .

Correct Approach:

1. p is prime and $p \equiv 1 \pmod{4}$, so $(-a)^2 = a^2$ modulo p .
2. The multiplicative order of $-a$ modulo p is the same as the multiplicative order of a modulo p , because:
 - $(-a)^k \equiv 1 \pmod{p}$ iff $a^k \equiv 1 \pmod{p}$ for any k , since $(-a)^k = (-1)^k * a^k$, and p is odd ($p \equiv 1 \pmod{4}$).
 - Alternatively, observe that $\text{orderOf } (-a) = \text{orderOf } a$ because $p \equiv 1 \pmod{4}$ and a is a primitive root.
 - In particular, $\text{orderOf } (-a) = p - 1 = \text{orderOf } a$, because $(-a)^{p-1} = a^{p-1} \equiv 1 \pmod{p}$ (Fermat's Little Theorem).

But we need a more precise argument:

1. If a is a primitive root, then the multiplicative order of a is $p - 1 = 4t$.
2. The multiplicative order of $-a$ is the same as the multiplicative order of a because $p \equiv 1 \pmod{4}$ and a is a primitive root.
 - To see this, note that $(-a)^2 = a^2$, so $\text{orderOf } (-a)$ divides $\text{orderOf } a^2$ (which is $2t$).
 - But $\text{orderOf } a^2$ divides $\text{orderOf } a = 4t$ (since $a^{2t} = (a^2)^t \equiv 1 \pmod{p}$).
 - Also, $\text{orderOf } (-a)$ divides $p - 1 = 4t$ because $(-a)^{4t} = ((-a)^2)^{2t} = (a^2)^{2t} \equiv 1 \pmod{p}$.
 - Since $\text{orderOf } (-a)$ divides $4t$ and $4t$ is the order of a , and $\text{orderOf } (-a)$ divides $\text{orderOf } a^2$ (which is $2t$), and $\text{orderOf } a^2$ divides $\text{orderOf } a$ (which is $4t$), the only common divisors are 1 and $2t$.

- But `orderOf a = 4*t` and `orderOf a^2 = 2*t`.
- Alternatively, because `a` is a primitive root, `a^k ≡ 1 mod p` iff `4*t` divides `k`.
- Similarly, `(-a)^k ≡ 1 mod p` iff `4*t` divides `k`, because `(-a)^k = a^k` (since `p ≡ 1 mod 4`).
- Therefore, the order of `-a` is the same as the order of `a`.

But we can simplify this:

- The order of `-a` is the same as the order of `a` because `p ≡ 1 mod 4` and `a` is a primitive root.

Thus, we have:

1. If `a` is a primitive root, then `-a` is also a primitive root.
2. Similarly, if `-a` is a primitive root, then `a` is also a primitive root.

Step 1: Abstract Plan

1. **Assume `a` is a primitive root modulo `p`:**
 - The multiplicative order of `a` is `p - 1 = 4*t`.
 - We need to show that the multiplicative order of `-a` is `4*t`.
 - Use the fact that `(-a)^k ≡ 1 mod p` iff `a^k ≡ 1 mod p`, because `p ≡ 1 mod 4`.
 - Thus, the order of `-a` is the same as the order of `a`.
 - Therefore, `-a` is also a primitive root.
2. **Assume `-a` is a primitive root modulo `p`:**
 - Similarly, `a` is a primitive root.
 - The reasoning is symmetric.

Step 2: Lean 4 `have` Statements

```

lean4
theorem exercise_4_4 {p t: ℕ} (hp0 : p.Prime) (hp1 : p = 4*t + 1)
  (a : ZMod p) :
  IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
  have h_main : IsPrimitiveRoot a p → IsPrimitiveRoot (-a) p := by
    sorry
  have h_converse : IsPrimitiveRoot (-a) p → IsPrimitiveRoot a p := by
    sorry
  have h_final : IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
    sorry
  exact h_final

```

Explanation:

1. `h_main`: If `a` is a primitive root, then `-a` is also a primitive root.
2. `h_converse`: If `-a` is a primitive root, then `a` is also a primitive root.
3. `h_final`: Combine the two directions to get the full biconditional.

Complete Lean 4 Proof

```

lean4
theorem exercise_4_4 {p t: ℕ} (hp0 : p.Prime) (hp1 : p = 4*t + 1)
  (a : ZMod p) :
  IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
  have h_main : IsPrimitiveRoot a p → IsPrimitiveRoot (-a) p := by
    intro h
    haveI := Fact.mk hp0
    have h1 := h.isUnit
    have h2 := h.pow_eq_one
    have h3 := h.pow_ne_one_of_pos_of_lt (by decide : (0 : ℕ) < 1) (by
      simp_all [hp1, Nat.Prime.one_lt]
    <|> omega)

```

```

have h4 := h.pow_ne_one_of_pos_of_lt (by decide : (0 : ℕ) < 2) (by
  simp_all [hp1, Nat.Prime.one_lt]
  <|> omega)
simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
<|>
(try omega) <|>
(try simp_all [IsPrimitiveRoot, pow_mul, pow_add, pow_one, pow_two,
  ↪ ZMod.nat_cast_self]) <|>
(try omega) <|>
(try
  {
    have h5 := h.pow_eq_one
    have h6 := h.isUnit
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <|>
    aesop
  }) <|>
(try
  {
    have h5 := h.pow_eq_one
    have h6 := h.isUnit
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <|>
    aesop
  })
<|>
(try
  {
    aesop
  })
<|>
(try
  {
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <|>
    aesop
  })
<|>
(try
  {
    omega
  })
<|>
(try
  {
    aesop
  })
<|>
(try
  {
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <|>
    omega
  })

have h_converse : IsPrimitiveRoot (-a) p → IsPrimitiveRoot a p := by
intro h
haveI := Fact.mk hp0
have h1 := h.isUnit
have h2 := h.pow_eq_one
have h3 := h.pow_ne_one_of_pos_of_lt (by decide : (0 : ℕ) < 1) (by
  simp_all [hp1, Nat.Prime.one_lt]
  <|> omega)
have h4 := h.pow_ne_one_of_pos_of_lt (by decide : (0 : ℕ) < 2) (by
  simp_all [hp1, Nat.Prime.one_lt]

```

```

    <.> omega)
simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
<.>
(tr try omega) <.>
(tr try simp_all [IsPrimitiveRoot, pow_mul, pow_add, pow_one, pow_two,
  ↪ ZMod.nat_cast_self]) <.>
(tr try omega) <.>
(tr try
  {
    have h5 := h.pow_eq_one
    have h6 := h.isUnit
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <.>
    aesop
  }) <.>
(tr try
  {
    have h5 := h.pow_eq_one
    have h6 := h.isUnit
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <.>
    aesop
  })
<.>
(tr try
  {
    aesop
  })
<.>
(tr try
  {
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <.>
    aesop
  })
<.>
(tr try
  {
    omega
  })
<.>
(tr try
  {
    aesop
  })
<.>
(tr try
  {
    simp_all [ZMod.nat_cast_self, pow_one, pow_two, neg_neg]
    <.>
    omega
  })
<.>
aesop

have h_final : IsPrimitiveRoot a p ↔ IsPrimitiveRoot (-a) p := by
  constructor
  · exact h_main
  · exact h_converse

exact h_final

```

D.5.2 CASE 2:

In another ProofNet example, **GAR** again succeeds where the base prover fails. This instance more clearly illustrates how **GAR** mitigates useless simplifications. It uses NLS to reach a comprehensive proof plan and effectively conduct the proof in a relatively simple manner. However, the base model conducts a quicker search but a prolonged proof. It expends effort on low-yield algebra and yields lengthy, fragile derivations. These observations highlight that the implicit curriculum in **GAR** fosters more efficient reasoning paths and greater verification robustness.

GAR-trained model generated result:

```
<|begin_of_sentence|><|User|>Complete the following Lean 4 code:

```lean4
import Mathlib
import Aesop

set_option maxHeartbeats 0

open BigOperators Real Nat Topology Rat

/-- Prove that $x^4+4x^3+6x^2+2x+1$ is irreducible in $\mathbb{Z}[x]$. -/
theorem exercise_9_4_2c : Irreducible
 (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial ℤ) := by
 ...
```

Before producing the Lean 4 code to formally prove the given theorem,  
 ↪ provide a detailed proof plan outlining the main proof steps and  
 ↪ strategies.

The plan should highlight key ideas, intermediate lemmas, and proof  
 ↪ structures that will guide the construction of the final formal  
 ↪ proof.

<|Assistant|>### Detailed Proof

**\*\*Problem:\*\*** Prove that  $\mathbb{P}(x) = x^4 + 4x^3 + 6x^2 + 2x + 1$  is  
 ↪ irreducible in  $\mathbb{Z}[x]$ .

**\*\*Approach:\*\***

We will use the Rabinowitz Criterion (also known as the "Rabinowitz  
 ↪ irreducibility test" or "Rabinowitz's Criterion") to prove that  $\mathbb{P}(x)$   
 ↪ is irreducible over  $\mathbb{Z}[x]$ .

**\*\*Rabinowitz Criterion:\*\***

Let  $\mathbb{P}(x) = x^n + a_{n-1}x^{n-1} + \dots + a_0$  be a polynomial in  
 ↪  $\mathbb{Z}[x]$ . Suppose that there exists a prime  $p$  such  
 ↪ that:

1.  $p$  divides each  $a_i$  for  $0 \leq i \leq n-1$ .
2.  $p^2$  does not divide  $a_0$ .

Then  $\mathbb{P}(x)$  is irreducible over  $\mathbb{Z}[x]$ .

**\*\*Application to  $\mathbb{P}(x)$ :\*\***

Let  $\mathbb{P}(x) = x^4 + 4x^3 + 6x^2 + 2x + 1$ .

1. The constant term is  $a_0 = 1$ .
2. The primes dividing all coefficients  $(a_0, a_1, a_2, a_3)$  are the  
 ↪ primes dividing  $\gcd(1, 2, 6, 4) = 1$ . So, no primes satisfy  
 ↪ the first condition.  
 - Hmm, this is incorrect: actually, all coefficients  $(a_0 = 1)$ ,  $(a_1 = 2)$ ,  $(a_2 = 6)$ ,  $(a_3 = 4)$  are divisible by  $1$ ,  
 ↪ but the second condition is about  $a_0 = 1$ , not all other  
 ↪ coefficients. This suggests that perhaps we need a different  
 ↪ approach.

But the Rabinowitz Criterion can still be applied:

Consider  $(p = 2)$ .

- $(2)$  divides  $(a_0 = 1)$  (False).
- $(2)$  divides  $(a_1 = 2)$  (True).
- $(2)$  divides  $(a_2 = 6)$  (True).
- $(2)$  divides  $(a_3 = 4)$  (True).

But  $(p^2 = 4)$  divides  $(a_0 = 1)$  (False).

This does not work. So, we need another prime.

Consider  $(p = 3)$ :

- $(3)$  divides  $(a_0 = 1)$  (No).
- $(3)$  divides  $(a_1 = 2)$  (No).

$(p)$  does not divide any of the coefficients, so this is invalid.

Consider  $(p = 5)$ :

- $(5)$  divides  $(a_0 = 1)$  (No).
- $(5)$  divides  $(a_1 = 2)$  (No).

$(p)$  does not divide any of the coefficients, so this is invalid.

### **\*\*Alternative Approach:\*\***

Let's consider the polynomial  $(Q(x) = x^4 + 4x^3 + 6x^2 + 2x + 1)$ . We  $\hookrightarrow$  can attempt to factor it or check for possible roots.

But  $(Q(0) = 1)$ ,  $(Q(1) = 1 + 4 + 6 + 2 + 1 = 14)$ , and  $(Q(-1) = 1 - 4 + 6 - 2 + 1 = 2)$ . No simple rational roots.

Alternatively, perhaps we can factor  $(Q(x))$  into quadratics.

Assume  $(Q(x) = (x^2 + ax + b)(x^2 + cx + d))$ .

Expanding gives:

$$(x^4 + (a+c)x^3 + (ac+b+d)x^2 + (ad+bc)x + bd = Q(x)).$$

Thus, we get the system:

1.  $(a + c = 4)$
2.  $(ac + b + d = 6)$
3.  $(ad + bc = 2)$
4.  $(bd = 1)$

From  $(bd = 1)$ , since  $(b, d \in \mathbb{Z})$ , we have the following  $\hookrightarrow$  cases:

1.  $(b = 1, d = 1)$ , or
2.  $(b = -1, d = -1)$ .

### **\*\*Case 1: $(b = d = 1)$ \*\***

From (1):  $(a + c = 4)$

From (2):  $(ac + 2 = 6 \implies ac = 4)$

From (3):  $(a + c = 4)$ ,  $(ac = 4)$

But  $(a, c)$  are roots of  $(t^2 - 4t + 4 = 0)$ , i.e.,  $(t = 2)$ . So  $\hookrightarrow (a = c = 2)$ .

This satisfies all equations:

- $(a + c = 4)$
- $(ac = 4)$
- $(ad + bc = 2 \cdot 1 + 2 \cdot 1 = 4 \neq 2)$  (Does not hold).

Wait, this is incorrect. The problem is in (3):

$(ad + bc = a \cdot 1 + c \cdot 1 = (a + c) = 4 \neq 2)$ , which  $\hookrightarrow$  contradicts the third condition.

Hence, this case is invalid.

### **\*\*Case 2: $(b = d = -1)$ \*\***

From (1):  $(a + c = 4)$

From (2):  $\backslash( ac - 2 = 6 \implies ac = 8 \backslash)$   
 From (3):  $\backslash( -a + -c = 2 \implies a + c = -2 \backslash)$ , but this contradicts  $\backslash( \leftrightarrow a + c = 4 \backslash)$ .

This case is also invalid.

**\*\*Another Approach: Eisenstein's Criterion\*\***

This polynomial is not directly suited for Eisenstein's Criterion,  
 $\leftrightarrow$  because replacing  $\backslash( x \backslash)$  with  $\backslash( x + 1 \backslash)$  gives:  
 $\backslash[ (x + 1)^4 + 4(x + 1)^3 + 6(x + 1)^2 + 2(x + 1) + 1 \backslash]$   
 $\backslash[ = x^4 + 4x^3 + 6x^2 + 4x + 1 + 4x^3 + 12x^2 + 12x + 4 + 6x^2 + 12x + 6$   
 $\leftrightarrow + 2x + 2 + 1 \backslash]$   
 $\backslash[ = x^4 + (4x^3 + 4x^3) + (6x^2 + 12x^2 + 6x^2) + (4x + 12x + 12x + 2x)$   
 $\leftrightarrow + (1 + 4 + 6 + 2 + 1) \backslash]$   
 $\backslash[ = x^4 + 8x^3 + 24x^2 + 30x + 14 \backslash]$   
 But we can try another substitution.

Alternatively, observe that:  
 $\backslash[ Q(x) = x^4 + 4x^3 + 6x^2 + 2x + 1 \backslash]$   
 $\backslash[ = (x^2 + 2x)^2 + 2x^2 + 2x + 1 \backslash]$   
 But this seems not helpful.

However, a better approach is to note that:  
 $\backslash[ Q(x) = (x^2 + 2x + 1)^2 - (2x^2 + 1) \backslash]$   
 But:  
 $\backslash[ Q(x) = x^4 + 4x^3 + 6x^2 + 2x + 1 \backslash]$   
 But:  
 $\backslash[ (x^2 + 2x + 1)^2 = x^4 + 4x^3 + 6x^2 + 4x + 1 \backslash]$   
 $\backslash[ Q(x) = (x^2 + 2x + 1)^2 - (2x + 2) \backslash]$   
 But  $\backslash( (x^2 + 2x + 1)^2 - (2x + 2) = x^4 + 4x^3 + 6x^2 + 4x + 1 - 2x - 2$   
 $\leftrightarrow = x^4 + 4x^3 + 6x^2 + 2x - 1 \neq Q(x) \backslash)$ .

This substitution is incorrect.

**\*\*Conclusion from the failed attempts:\*\***

It seems difficult to factor  $\backslash( Q(x) \backslash)$  in  $\backslash(\mathbb{Z}[x]\backslash)$  by simple  
 $\leftrightarrow$  methods.

But we can try a more general approach using the Cohn's Irreducibility  
 $\leftrightarrow$  Criterion or Ritt's Criterion, but these are not straightforward.

However, we can use the Rational Root Test (or simply checking small  
 $\leftrightarrow$  primes):

Check for  $\backslash( x = \pm 1, \pm p \backslash)$  for primes  $\backslash( p \backslash)$ :

- $\backslash( Q(1) = 1 + 4 + 6 + 2 + 1 = 14 \backslash)$
- $\backslash( Q(-1) = 1 - 4 + 6 - 2 + 1 = 2 \backslash)$
- $\backslash( Q(2) = 16 + 32 + 24 + 4 + 1 = 77 \backslash)$
- $\backslash( Q(-2) = 16 - 32 + 24 - 4 + 1 = 5 \backslash)$

Since  $\backslash( Q(x) \backslash)$  has no rational roots, it is irreducible over  
 $\leftrightarrow \backslash(\mathbb{Z}\backslash)$ .

But this seems incorrect, as  $\backslash( Q(x) \backslash)$  is reducible for  $\backslash( x = 0 \backslash)$  and  
 $\leftrightarrow \backslash( x = -1 \backslash)$ , but we are not using this approach.

Alternatively, we can use the **\*\*Schinzel's criterion\*\*** or **\*\*Rabinowitz**  
 $\leftrightarrow$  **Criterion\*\*** again:

Consider  $\backslash( Q(x + 1) = (x + 1)^4 + 4(x + 1)^3 + 6(x + 1)^2 + 2(x + 1) + 1$   
 $\leftrightarrow \backslash)$   
 $\backslash[ = x^4 + 4x^3 + 6x^2 + 4x + 1 + 4x^3 + 12x^2 + 12x + 4 + 6x^2 + 12x + 6$   
 $\leftrightarrow + 2x + 2 + 1 \backslash]$   
 $\backslash[ = x^4 + (4x^3 + 4x^3) + (6x^2 + 12x^2 + 6x^2) + (4x + 12x + 12x + 2x)$   
 $\leftrightarrow + (1 + 4 + 6 + 2 + 1) \backslash]$   
 $\backslash[ = x^4 + 8x^3 + 24x^2 + 30x + 14 \backslash]$

Hmm, no immediate factorization here either.

**\*\*Final Attempt:\*\***

We use the fact that if  $Q(x)$  is irreducible over  $\mathbb{Z}$ ,  
 $\hookrightarrow$  then it is irreducible over  $\mathbb{Q}$ .

Alternatively, we can directly check that  $Q(x)$  has no factors in  
 $\hookrightarrow \mathbb{Z}[x]$  other than  $Q(x)$  itself and constants.

We can perform a complete check:

Suppose  $Q(x) = A(x) B(x)$  is a factorization into non-constant  
 $\hookrightarrow$  polynomials.

We can write:

$$A(x) = x^4 + a_1 x^3 + a_2 x^2 + a_3 x + a_4$$

$$B(x) = 1$$

But  $B(x)$  is a unit and  $\deg(B(x)) = 0$ , so this is trivial.

Alternatively, assume  $A(x)$  has degree 1 or 2, etc.

But this is tedious and unlikely to work.

Given that the problem is from an older source, we can infer that the  
 $\hookrightarrow$  polynomial  $Q(x)$  is irreducible.

But in Lean, we can proceed by proving that  $Q(x)$  has no non-trivial  
 $\hookrightarrow$  factors, by checking that for every pair of integers  $(a, b)$   
 $\hookrightarrow$  such that  $b \neq \pm 1$ , we cannot write  $Q(x)$  as a product  
 $\hookrightarrow$  of two non-constant polynomials.

Alternatively, we can use Lean's libraries to prove irreducibility by  
 $\hookrightarrow$  contradiction:

**\*\*Proof Sketch:\*\***

1. Assume  $Q(x) = A(x) B(x)$ .
2. Look at the constant term.  
 The constant term is  $\pm 1$ , so possible constant factors are  $\pm 1$ .
3. So either  $A(x)$  or  $B(x)$  is a constant, say  $A(x) = \pm 1$ .
4. But then  $Q(x)$  would have to be constant, which is false.  
 So, contradiction.

But Lean may not directly have lemmas for this.

A better approach is to use that for a monic polynomial with integer  
 $\hookrightarrow$  coefficients and  $\pm 1$  constant term, it is irreducible if it  
 $\hookrightarrow$  has no non-trivial factors.

But in Lean, we can prove that no non-trivial factors exist.

However, checking all possible factors is impractical.

Alternatively, we can use the fact that  $Q(x)$  is a cyclotomic-like  
 $\hookrightarrow$  polynomial or a Schur polynomial.

But since this seems too involved, we can use Lean's **\*\*irreducibility testing\*\***  
 $\hookrightarrow$  (but that's not part of standard libraries).

**\*\*Alternative Step:\*\***

We can instead show that  $Q(x)$  is not a perfect square in  
 $\hookrightarrow \mathbb{Z}[x]$ .

But this is non-trivial.

Alternatively, we can use **\*\*Mignotte's bound\*\*** to check that no  
 $\hookrightarrow$  non-trivial factors exist.

But given that  $Q(x)$  is not a known factorizable polynomial, we can  
 $\hookrightarrow$  simply use the fact that it is not reducible and is of degree 4  
 $\hookrightarrow$ , so it must be irreducible if it has no roots in  $\mathbb{Z}$ .

Since  $Q(x)$  is strictly increasing for  $x \geq 0$  and  $Q(-1) = 2$ ,  
 $Q(0) = 1$ , and  $Q(1) = 14$ , there are no integer  
 $\hookrightarrow$  roots.

But Lean can confirm this by checking values.

But instead, we can use:

**\*\*Lemmas:\*\***

1. If  $Q(x)$  is reducible in  $\mathbb{Z}[x]$ , then it has a  
 $\hookrightarrow$  linear or quadratic factor.
2. A reducible polynomial must have a factor with  $\deg \leq \lfloor \frac{\deg(Q)}{2} \rfloor$ .  
 $\hookrightarrow$
3. The possible quadratics are  $x^2 + ax + 1$  or similar, but none  
 $\hookrightarrow$  divide  $Q(x)$ .

But checking this is tedious.

Instead, we can appeal to a known fact that  $x^4 + x^3 + x^2 + x + 1$   
 $\hookrightarrow$  is irreducible, and substituting  $x + 1$  gives:

$$\begin{aligned} & [(x + 1)^4 + 4(x + 1)^3 + 6(x + 1)^2 + 2(x + 1) + 1] \\ & = x^4 + 4x^3 + 6x^2 + 4x + 1 + 4x^3 + 12x^2 + 12x + 4 + 6x^2 + 12x + 6 \\ & \quad + 2x + 2 + 1 \\ & = x^4 + 8x^3 + 24x^2 + 30x + 14 \end{aligned}$$

This is the same as above.

Alternatively, check that  $x^4 + 4x^3 + 6x^2 + 2x + 1$  is irreducible  
 $\hookrightarrow$  using:

**\*\*Useful Lemma:\*\*** The polynomial  $x^4 + ax^3 + bx^2 + cx + 1$  is  
 $\hookrightarrow$  irreducible in  $\mathbb{Z}[x]$  if  $a^2 < 4b$  or  $b^2 < 4c$   
 $\hookrightarrow$ , etc.

But this is not directly helpful.

**\*\*Lean-Applicable Approach:\*\***

To prove irreducibility in Lean, we can directly use:

- The **\*\*reduction modulo primes\*\*** approach to eliminate possible  
 $\hookrightarrow$  factorizations.

Alternatively, we can use the fact that  $Q(x)$  is a cyclotomic-like  
 $\hookrightarrow$  polynomial or a Cohn polynomial.

But given that the polynomial is of degree 4 and no obvious  
 $\hookrightarrow$  factorization exists, we can proceed by contradiction.

**### Step-by-Step Abstract Plan**

1. **\*\*Assume Reducibility\*\***: Suppose  $Q(x)$  is reducible in  $\mathbb{Z}[x]$ . Then,  $Q(x) = A(x)B(x)$  with  $\deg(A(x)) \geq 1$ ,  $\deg(B(x)) \geq 1$ , and  $\deg(A(x)) + \deg(B(x)) = 4$ .
2. **\*\*Analyze Constant Terms\*\***: Since  $Q(0) = 1$  and  $Q(x)$  has no rational roots, no constant factors can exist unless one of  $A(x)$  or  $B(x)$  is  $\pm 1$ .

3. **\*\*Check for Linear Factors\*\***: If  $\deg(A(x)) = 1$ , say  $A(x) = x + a$ , then  $Q(x) = (x + a)B(x)$ . Expand to get coefficients and contradict.

4. **\*\*Check for Quadratic Factors\*\***: If  $\deg(A(x)) = 2$ , say  $A(x) = x^2 + ax + b$ , analyze to find contradictions.

5. **\*\*Contradiction\*\***: Since all possible factorizations lead to contradictions,  $Q(x)$  must be irreducible.

### ### Lean 4 Abstract Plan with `have` Statements

```

```lean4
theorem exercise_9_4_2c : Irreducible
  (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial ℤ) := by
  have h_main : Irreducible (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial
    ↪ ℤ) := by
    sorry
  sorry
```

```

### ### Complete Lean 4 Proof

```

```lean4
theorem exercise_9_4_2c : Irreducible
  (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial ℤ) := by
  have h_main : Irreducible (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial
    ↪ ℤ) := by
    apply?
  apply h_main
```

```

### Base model generated result:

<|begin\_of\_sentence|><|User|>Complete the following Lean 4 code:

```

```lean4
import Mathlib
import Aesop

set_option maxHeartbeats 0

open BigOperators Real Nat Topology Rat

/-- Prove that  $x^4+4x^3+6x^2+2x+1$  is irreducible in  $\mathbb{Z}[x]$ . -/
theorem exercise_9_4_2c : Irreducible
  (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial ℤ) := by
```

```

Before producing the Lean 4 code to formally prove the given theorem,  
 ↪ provide a detailed proof plan outlining the main proof steps and  
 ↪ strategies.

The plan should highlight key ideas, intermediate lemmas, and proof  
 ↪ structures that will guide the construction of the final formal  
 ↪ proof.

<|Assistant|>### Detailed Proof and Analysis

We need to prove that the polynomial  $P(x) = x^4 + 4x^3 + 6x^2 + 2x + 1$   
 ↪ is irreducible over  $\mathbb{Z}[x]$ .

### #### Step 1: Understand the Polynomial

First, observe that  $(P(x))$  is palindromic. A palindromic polynomial  $\leftrightarrow (P(x))$  can be factored as  $(P(x) = (x^2 + a x + b)(x^2 + c x + d))$  if and only if  $(P(x))$  can be written as a product of two quadratic polynomials.

#### #### Step 2: Try to Factor $(P(x))$

Assume that  $(P(x))$  factors into two non-constant polynomials with integer coefficients:

$$(P(x) = (x^2 + a x + b)(x^2 + c x + d)).$$

Expanding the right-hand side gives:

$$(x^4 + (a + c)x^3 + (ac + b + d)x^2 + (ad + bc)x + bd).$$

By comparing coefficients, we get the system:

1.  $(a + c = 4)$ ,
2.  $(ac + b + d = 6)$ ,
3.  $(ad + bc = 2)$ ,
4.  $(bd = 1)$ .

#### #### Step 3: Solve the System

From  $(bd = 1)$ , the possible integer pairs  $((b, d))$  are:

$$((1, 1), (-1, -1)).$$

**\*\*Case 1:\*\***  $(b = 1)$ ,  $(d = 1)$ .

Substitute into the other equations:

1.  $(a + c = 4)$ ,
2.  $(ac + 1 + 1 = 6 \rightarrow ac = 4)$ ,
3.  $(a + c = 4)$ ,
4.  $(ad + bc = a + c = 4)$ , but this is already satisfied.

But from  $(a + c = 4)$  and  $(ac = 4)$ , the quadratic equation for  $(a)$  is:

$$(t^2 - 4t + 4 = 0 \rightarrow (t - 2)^2 = 0 \rightarrow t = 2).$$

Thus  $(a = c = 2)$ .

Check the last equation:  $(ad + bc = 2 \cdot 1 + 2 \cdot 1 = 4 \neq 2)$

$\leftrightarrow$ ). This is a contradiction, so this case is invalid.

**\*\*Case 2:\*\***  $(b = -1)$ ,  $(d = -1)$ .

Substitute into the other equations:

1.  $(a + c = 4)$ ,
2.  $(ac - 1 - 1 = 6 \rightarrow ac = 8)$ ,
3.  $(ad + bc = -a - c = 2)$ ,
4.  $(bd = 1)$ .

But from  $(a + c = 4)$  and  $(ac = 8)$ , the quadratic equation for  $(a)$  is:

$$(t^2 - 4t + 8 = 0),$$

which has discriminant  $(D = 16 - 32 = -16 < 0)$ . Thus, there are no

$\leftrightarrow$  real (and hence no integer) roots, so this case is invalid.

#### #### Step 4: Check Other Possible Factorizations

The above exhaustive search shows that  $(P(x))$  does not factor into  $\leftrightarrow$  two non-constant polynomials with integer coefficients. Thus,  $(P(x))$  is irreducible.

#### #### Verification of the Irreducibility

We have checked all possible factorizations of  $(P(x))$  into quadratics

$\leftrightarrow$  and found that none satisfy the coefficient conditions. Therefore,  $(P(x))$

$\leftrightarrow (P(x))$  is irreducible.

#### ### Abstract Plan

1. **\*\*Assume Factorization\*\***: Suppose  $(P(x) = (x^2 + a x + b)(x^2 + c x + d))$  for integers  $((a, b, c, d))$ .

2. **\*\*Compare Coefficients\*\***: Derive the system of equations for  $(a, b, c, d)$ .
  - $(a + c = 4)$ ,
  - $(ac + b + d = 6)$ ,
  - $(ad + bc = 2)$ ,
  - $(bd = 1)$ .
3. **\*\*Solve  $(bd = 1)$ \*\***: Consider the pairs  $(b, d)$  as  $((1, 1))$  and  $((-1, -1))$ .
4. **\*\*Case  $(b = 1), (d = 1)$ \*\***:
  - Derive  $(a + c = 4)$  and  $(ac = 4)$ .
  - Solve for  $(a)$  and  $(c)$  to find  $(a = c = 2)$ .
  - But this leads to  $(ad + bc = 4 \neq 2)$ , a contradiction.
5. **\*\*Case  $(b = -1), (d = -1)$ \*\***:
  - Derive  $(a + c = 4)$  and  $(ac = 8)$ .
  - Show that no real roots exist for  $(a)$ , so this case is invalid.
6. **\*\*Conclusion\*\***: No valid factorization exists, so  $(P(x))$  is irreducible.

### ### Lean 4 `have` Statements

```

```lean4
theorem exercise_9_4_2c : Irreducible
  (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial ℤ) := by
  have h_main : Irreducible (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial
    ↪ ℤ) := by sorry
  sorry
```

```

### ### Complete Lean 4 Proof

```

```lean4
theorem exercise_9_4_2c : Irreducible
  (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial ℤ) := by
  have h_main : Irreducible (X^4 + 4*X^3 + 6*X^2 + 2*X + 1 : Polynomial
    ↪ ℤ) := by
    -- Assume the polynomial factors into two non-constant polynomials
    -- with integer coefficients
    refine' (prime_of_irreducible (show Prime (X^4 + 4*X^3 + 6*X^2 + 2*X
      ↪ + 1 : Polynomial ℤ) by
      -- Use the fact that the polynomial is prime to show it is
      ↪ irreducible
      apply prime_C_mul_X_pow_add_C_of_prime
      <.> norm_num
      <.>
      decide))
    <.>
    (try norm_num) <.>
    (try decide) <.>
    (try
      {
        intro h
        -- If the polynomial factors, it must factor into quadratics
        rcases h with ⟨p, q, hp, hq, hpq⟩
        -- Analyze the possible factorizations
        have h1 := congr_arg (fun p => eval 0 p) hpq
        have h2 := congr_arg (fun p => eval 1 p) hpq
        have h3 := congr_arg (fun p => eval (-1) p) hpq
        have h4 := congr_arg (fun p => eval 2 p) hpq
        have h5 := congr_arg (fun p => eval (-2) p) hpq
        simp at h1 h2 h3 h4 h5
        -- Check the possible evaluations to find a contradiction
        norm_num [eval_add, eval_mul, eval_pow, eval_C, eval_X, eval_pow,
          ↪ eval_mul, eval_add, eval_sub] at h1 h2 h3 h4 h5
        <.>
        (try { contradiction }) <.>
      }
    )
  )
```

```

```

(try {
 have h6 := p.sub_dvd_eval_sub 0 1
 have h7 := q.sub_dvd_eval_sub 0 1
 have h8 := p.sub_dvd_eval_sub 0 (-1)
 have h9 := q.sub_dvd_eval_sub 0 (-1)
 have h10 := p.sub_dvd_eval_sub 0 2
 have h11 := q.sub_dvd_eval_sub 0 2
 have h12 := p.sub_dvd_eval_sub 0 (-2)
 have h13 := q.sub_dvd_eval_sub 0 (-2)
 simp [eval_add, eval_mul, eval_pow, eval_C, eval_X] at h6 h7 h8
 ↪ h9 h10 h11 h12 h13
<.>
(try omega) <.>
(try {
 have h14 : p.natDegree ≤ 2 := by
 contrapose! h6
 simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
 <.>
 (try omega) <.>
 (try nlinarith)
 have h15 : q.natDegree ≤ 2 := by
 contrapose! h7
 simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
 <.>
 (try omega) <.>
 (try nlinarith)
 have h16 : p.natDegree ≤ 2 := by
 contrapose! h8
 simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
 <.>
 (try omega) <.>
 (try nlinarith)
 have h17 : q.natDegree ≤ 2 := by
 contrapose! h9
 simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
 <.>
 (try omega) <.>
 (try nlinarith)
 have h18 : p.natDegree ≤ 2 := by
 contrapose! h10
 simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
 <.>
 (try omega) <.>
 (try nlinarith)
 have h19 : q.natDegree ≤ 2 := by
 contrapose! h11
 simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
 <.>
 (try omega) <.>
 (try nlinarith)
 have h20 : p.natDegree ≤ 2 := by
 contrapose! h12

```

```

simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
<;>
(tr try omega) <;>
(tr try nlinarith)
have h21 : q.natDegree ≤ 2 := by
 contrapose! h13
 simp_all [Polynomial.eval_add, Polynomial.eval_mul,
 ↪ Polynomial.eval_pow, Polynomial.eval_C,
 ↪ Polynomial.eval_X]
 <;>
 (tr try omega) <;>
 (tr try nlinarith)
-- Now we know the degree of p and q is at most 2
have h22 : p.natDegree ≤ 2 := by linarith
have h23 : q.natDegree ≤ 2 := by linarith
have h24 : p = C (p.coeff 0) ∨ p = C (p.coeff 0) * X ∨ p = C
 ↪ (p.coeff 0) * X ^ 2 := by
 have h25 := p.natDegree_le_of_degree_le h22
 have h26 := p.natDegree_le_of_degree_le h22
 have h27 := p.natDegree_le_of_degree_le h22
 cases' lt_or_eq_of_le (natDegree_le_of_degree_le h22) with
 ↪ h28 h28 <;>
 simp_all [Polynomial.natDegree_le_iff_degree_le,
 ↪ Polynomial.degree_eq_natDegree]
 <;>
 (tr try omega) <;>
 (tr try aesop) <;>
 (tr try {
 aesop
 }) <;>
 (tr try {
 aesop
 }) <;>
 (tr try {
 aesop
 })
have h25 : q = C (q.coeff 0) ∨ q = C (q.coeff 0) * X ∨ q = C
 ↪ (q.coeff 0) * X ^ 2 := by
 have h26 := q.natDegree_le_of_degree_le h23
 have h27 := q.natDegree_le_of_degree_le h23
 have h28 := q.natDegree_le_of_degree_le h23
 cases' lt_or_eq_of_le (natDegree_le_of_degree_le h23) with
 ↪ h29 h29 <;>
 simp_all [Polynomial.natDegree_le_iff_degree_le,
 ↪ Polynomial.degree_eq_natDegree]
 <;>
 (tr try omega) <;>
 (tr try aesop) <;>
 (tr try {
 aesop
 }) <;>
 (tr try {
 aesop
 }) <;>
 (tr try {
 aesop
 })
 })
rcases h24 with (h25 | h25 | h25) <;> rcases h25 with (h26 | h26
 ↪ | h26) <;> simp_all [Polynomial.eval_add,
 ↪ Polynomial.eval_mul, Polynomial.eval_pow,
 ↪ Polynomial.eval_C, Polynomial.eval_X]
<;>
(tr try omega) <;>

```

```
(try nlinarith) <.>
 (try {
 ring_nf at *
 <.>
 omega
 }) <.>
 (try {
 norm_num at *
 <.>
 aesop
 })
 })
 <.>
 aesop
exact h_main

```