

HARNESSING PROOF ASSISTANT FEEDBACK FOR REINFORCEMENT LEARNING AND MONTE-CARLO TREE SEARCH

Anonymous authors

Paper under double-blind review

ABSTRACT

Lean is an advanced proof assistant designed to facilitate formal theorem proving by providing a variety of interactive feedback. In this paper, we explore methodologies to leverage proof assistant feedback to augment the capabilities of large language models in constructing formal proofs. First, we deploy online reinforcement learning using Lean verification outcomes as the reward signal to improve the proof completion policy. This straightforward approach shows great promise in enhancing the model’s alignment with the formal verification system. In addition, we propose RMaxTS, a variant of Monte-Carlo tree search that employs an intrinsic-reward-driven exploration strategy to generate diverse proof paths. The tree structure is organized to represent the transitions of intermediate tactic states, extracted from the compilation messages given by Lean’s tactic mode. The intrinsic reward is constructed to incentivize the discovery of novel tactic states, which helps to mitigate the sparse-reward problem inherent in proof search. These techniques lead to a more efficient planning scheme for formal proof generation, achieving new state-of-the-art results on both miniF2F and ProofNet benchmarks.

1 INTRODUCTION

Recent advancements in large language models have significantly influenced mathematical reasoning and theorem proving in artificial intelligence. Despite notable progress in natural language domains, language models still encounter substantial challenges in formal theorem proving, *e.g.* using Lean (Moura & Ullrich, 2021) and Isabelle (Paulson, 1994), which requires rigorous derivations satisfying formal specifications of the verification system. Even advanced models like GPT-4 (OpenAI, 2023) struggle with complex formal proofs, underscoring the intricate nature of both the coding and the mathematics involved. A formal theorem proving model must not only grasp the syntax and semantics of formal systems like the Lean theorem prover but also align abstract mathematical reasoning with precise formal representation.

Language models in formal theorem proving typically employ two strategies: proof-step generation (Jiang et al., 2022a; Lample et al., 2022; Yang et al., 2023; Wu et al., 2024) and whole-proof generation (Zhao et al., 2023; Wang et al., 2023a). The proof-step generation approach is motivated by the interactive nature of Lean’s tactic mode, in which the compiler provides the access to the tactic state, *i.e.*, a structured representation summarizing the current status of the proof, including all the relevant information such as the local context of hypotheses and pending goals. Given the intermediate tactic state, the proof-step generation approach predicts each subsequent tactic and verifies it using the formal verifier to obtain updated information about the current tactic state. This interactive process often employs tree search techniques to compose valid proofs through several iterations of tactic generation (Polu & Sutskever, 2020). In contrast, the whole-proof generation approach treats the construction of formal proofs as a general code completion task. This branch of methods aims to generate the entire proof code based on the theorem statement and perform verification only at the end of the generation process. The simplicity of the whole-proof generation paradigm has been proven to offer high scalability (Xin et al., 2024) from the perspectives of both model training and inference deployment. In addition, the whole-proof generation model is trained to perform long-term planning for theorem proving, facilitating the integration and utilization of the model’s capabilities in natural language mathematical reasoning (Jiang et al., 2022b).

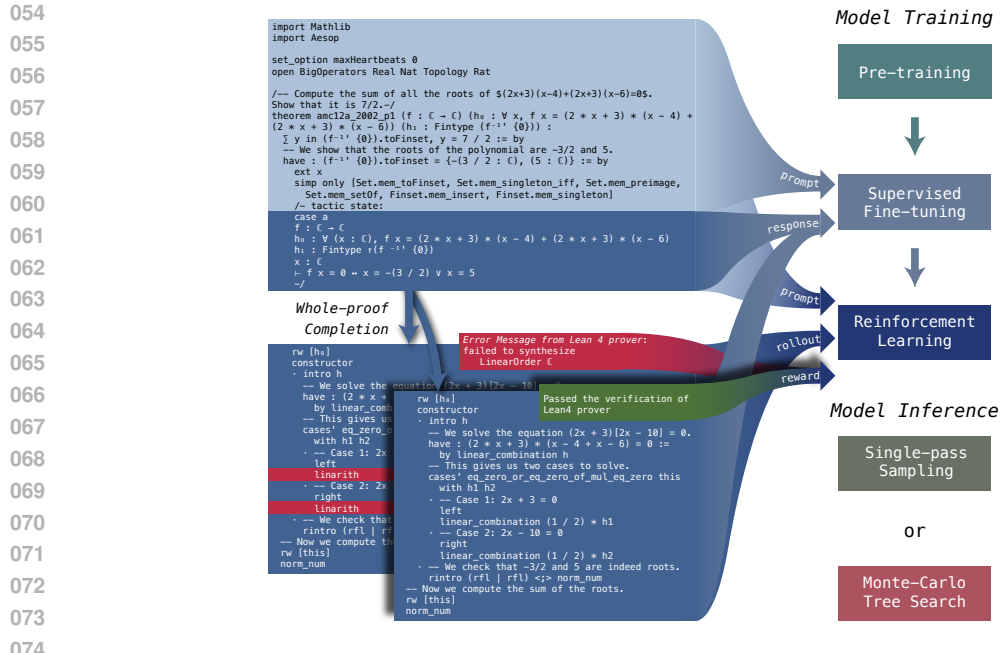


Figure 1: **Overall Framework of DS-Prover-V1.5.** During supervised fine-tuning, the model receives an incomplete theorem proof ending with a tactic state comment keyword. The model is trained to predict the content of this tactic state (auxiliary objective) and complete the subsequent proof steps (main objective). In the reinforcement learning stage, given an incomplete theorem proof and ground-truth tactic state from the Lean prover, we roll out the fine-tuned model to generate multiple proof candidates, which are then verified by the Lean prover. The verification results for these candidates are used as binary rewards to further optimize the model and enhance its alignment with the formal specifications of the verification system. For model inference, we decompose the generated proof into a series of tree nodes, appending intermediate tactic states extracted from the Lean prover, thereby establishing an interactive proof search paradigm.

In this paper, we present a unified approach that combines the strengths of both proof-step and whole-proof generation paradigms. We begin by training a whole-proof generation model, incorporating several auxiliary tasks to enhance its capabilities in mathematical reasoning and long-horizon planning, meanwhile empowering it to recognize information from Lean’s proof assistant feedback. The model is named DS-Prover-V1.5, as it builds upon the prior work of DeepSeek-Prover-V1 (Xin et al., 2024). We then employ a truncate-and-resume mechanism to decompose the whole-proof generation into a tactic-level proof search scheme. Figure 1 presents an illustration of our approach. The process begins with standard whole-proof generation, where the language model completes the proof code following the theorem statement prefix. The Lean assistant then verifies this code. If an error is detected, the code is truncated at the first error message, and any subsequent code is discarded. The successfully generated proof code is then used as a prompt for the generation of next proof segment. The latest tactic state from the Lean prover is appended at the end of the prompt as a comment block to provide intermediate guidance for the construction of long proofs. Notably, our method is not restricted to resuming from the last successfully applied tactic. We formalize the truncate-and-resume mechanism within the framework of Monte-Carlo tree search (MCTS; Couloum, 2006) in which the truncation points are scheduled by the tree search policy. In addition, we propose a novel reward-free exploration algorithm for MCTS to address the reward sparsity issue of proof search. We assign the tree search agent intrinsic motivation, *a.k.a.* curiosity (Schmidhuber, 2010), to extensively explore the tactic state space. These algorithmic modules extend the functionality of our whole-proof generation model to become a flexible tool for interactive theorem proving, which can effectively utilize the proof assistant feedback and generate diverse solution candidates. In experiments, we demonstrate substantial improvement of our proposed approach over baseline models, achieving new state-of-the-art results on the test set of the high school level miniF2F benchmark (63.5%) and the undergraduate level ProofNet benchmark (25.3%).

2 RELATED WORK

Reinforcement Learning for Theorem Proving. Numerous prior research efforts have explored modeling the interaction interface with the proof assistants as a Markov Decision Process (MDP), leveraging various reinforcement learning techniques, such as using policy gradient (Zombori et al., 2021; Crouse et al., 2021) and deep Q-learning (Fawzi et al., 2019), and involving a wide range of formal verification systems, including E (McKeown & Sutcliffe, 2023), Coq (Kusumoto et al., 2018), and HOL4 (Gauthier, 2020). A common choice for the reward signal is a binary indicator denoting whether the proof has been completed. Fawzi et al. (2019) designed a temporal-difference reward assignment according to problem structure of solving polynomial inequality. Aygün et al. (2022) generalized the idea of hindsight experience replay (Andrychowicz et al., 2017) from goal-reaching control to formal theorem proving, which enriches the reward supervision.

Tree Search for Theorem Proving. Integrating supervised models with search algorithms is a classical paradigm for automated theorem proving (Rawson & Reger, 2019; 2021; Zhang et al., 2024). For proof-step generation models, the most widely applied search strategy is best-first search (Yang et al., 2023), in which search branches are prioritized based on the cumulative log-likelihoods of the generated tactics. Lample et al. (2022) developed a specialized Monte-Carlo tree search algorithm tailored for the Lean theorem prover, in which subgoal branches are represented as hyperedges. The model training and tree search procedures are integrated similarly to the algorithmic framework of AlphaZero (Silver et al., 2018). Beyond the tactic-level tree abstraction, Wang et al. (2023b) investigated the effectiveness of using a proof-level value function in proof tree search, demonstrating that incorporating the entire proof as context is more effective than using a tactic-level state representation.

3 LEARNING TO UTILIZE PROOF ASSISTANT FEEDBACK

3.1 SUPERVISED FINE-TUNING

In this section, we explore the methodology and processes involved in the supervised fine-tuning (SFT) of DS-Prover-V1.5. Specifically, we incorporate intermediate tactic state information as an auxiliary prediction task to support the truncate-and-resume mechanism used in Monte-Carlo tree search. In addition, we augment the proof dataset from DeepSeek-Prover-V1 (Xin et al., 2024) by adding detailed explanatory comments. This enhancement aims to improve the alignment between natural language descriptions and Lean 4 code, thereby facilitating better formal mathematical reasoning. We refer to the resulting model as DS-Prover-V1.5-SFT. Details of data processing are described in Appendix A.2.

Prompt Augmentation with Tactic State Information. To implement the truncate-and-resume mechanism for Monte-Carlo Tree Search, we needed to extract tactic information from the code generated by the model. We enhanced the Lean REPL (Read-Eval-Print Loop; Leanprover Community, 2023) with data extraction tools from the LeanDojo (Yang et al., 2023) project. This allowed us to extract tactic information in triples, which include the position of each tactic, as well as the tactic states before and after its application. This information helps us identify the specific tactic code that triggers verification errors (used in the expansion step for tree search, see Section 4.2). For each tactic in a generated valid formal proof, we insert the tactic state returned by the verifier as a comment “/- tactic state: ... -/”. During training, we include all tokens following the leading prompt “/- tactic state: ” as responses to calculate the supervised fine-tuning loss, while the tokens before this comment is used as prompts and do not contribute to the training loss calculation, *i.e.*, we construct an auxiliary task for the prover model to predict the current tactic state.

Thought-augmented Proof Generation. Similar to Lean-STaR (Lin et al., 2024), which performs isolated chain-of-thought reasoning (Wei et al., 2022; Feng et al., 2023) before generating each proof step, we integrate this reasoning procedure directly as comments within the proof code. We use the DeepSeek-Coder V2 236B (Zhu et al., 2024) to enhance existing data in DeepSeek-Prover-V1 in two ways: first, by inserting a complete natural language solution at the beginning of the proof block, and second, by alternately inserting specific natural language steps for corresponding Lean tactics. Training the model with this data format enforces it to propose complete

162 mathematical reasoning at the beginning of the proof block and detailed step planning before each
 163 tactic. This approach successfully develops new behaviors, employing delicate mathematical think-
 164 ing to guide the generation of tactics. In the training data, two distinct guiding prompts are used
 165 to differentiate between the CoT (Chain of Thought) mode and the non-CoT mode for proof code
 166 completion. Examples of input and output in both modes can be found in Appendix C.

167
 168 **Discussion.** The primary purpose of implementing these data processing procedures during the
 169 SFT phase is to optimize the model’s performance for downstream inference-time strategy. When
 170 applying tree search for proof generation, we leverage the model’s ability to utilize proof assistant
 171 feedback while performing chain-of-thought reasoning. For an incomplete proof, we first append a
 172 comment block containing the ground-truth tactic state extracted from the Lean assistant, and then
 173 the model would perform chain-of-thought based on the full context information (see Figure 2). This
 174 procedure emulates the strategy employed by human experts when interacting with a formal proof
 175 assistant, combining both real-time feedback from the assistant and logical reasoning to iteratively
 176 refine and extend the proof.

177 3.2 REINFORCEMENT LEARNING FROM PROOF ASSISTANT FEEDBACK

178
 179 Reinforcement learning (RL) has been proven effective in enhancing the mathematical reasoning
 180 capabilities of supervised fine-tuned language models (Shao et al., 2024). To further advance DS-
 181 Prover-V1.5-SFT, we incorporate a reinforcement learning phase, resulting in the model DS-Prover-
 182 V1.5-RL. This phase leverages RL to enhance performance based on verification feedback from the
 183 Lean 4 prover. The specifics of this RL process are detailed below. Detailed training setting and
 184 hyper-parameters refer to Appendix A.3.

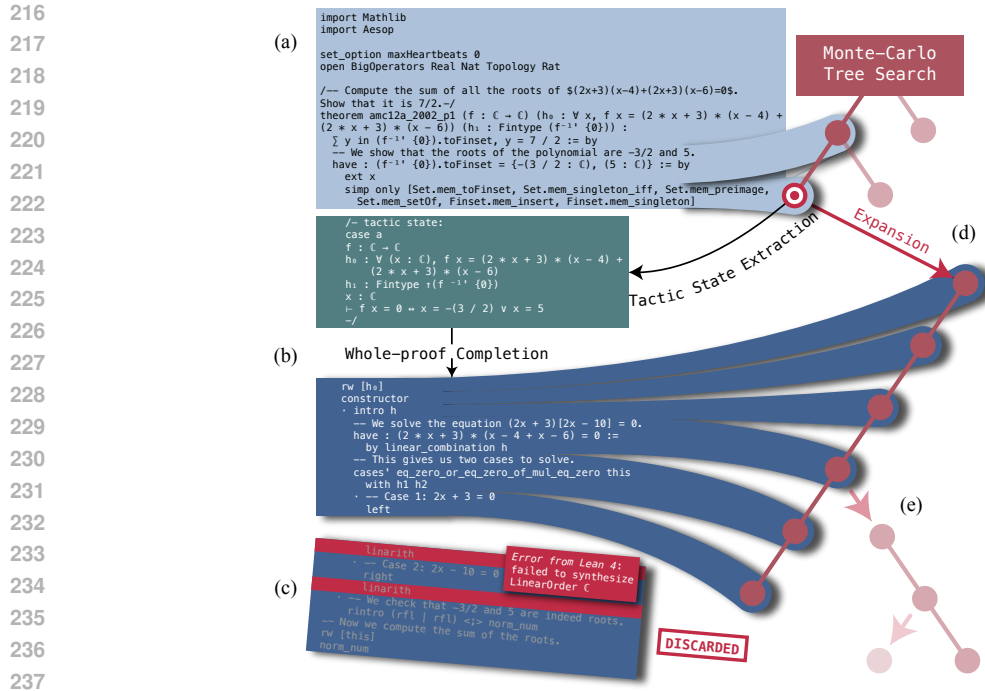
185
 186 **Reinforcement Learning Algorithm.** We employ Group Relative Policy Optimization (GRPO;
 187 Shao et al., 2024) as our RL algorithm, which has demonstrated superior effectiveness and efficiency
 188 compared to PPO (Schulman et al., 2017), primarily because it eliminates the necessity of training
 189 an additional critic model. Specifically, GRPO samples a group of candidate proofs for each theorem
 190 prompt and optimizes the model based on the relative rewards of the outputs within the group. To
 191 ensure that both correct and incorrect proofs are included in the rollout candidates, we select a
 192 subset of theorem statements with appropriate difficulty from the supervised fine-tuning dataset as
 193 training prompts. These selected prompts are chosen based on the rule that DS-Prover-V1.5-SFT
 194 achieves a moderate success rate in generating correct proofs over multiple attempts. After filtering,
 195 we retain approximately 4.5k unique theorem statements. Each theorem is prefixed with both CoT
 196 and non-CoT guiding prompts to enhance the model’s proof generation capabilities in both modes.
 197 The reward function is naturally given by the formal verification system, *i.e.*, each generated proof
 198 receives a reward of 1 if verified as correct, and 0 otherwise.

199
 200 **Discussion.** In Lean’s tactic mode, proofs are constructed through a sequence of tactics that trans-
 201 form the proof state. This sequential nature introduces the risk of compounding errors (Ross et al.,
 202 2011), where a single misinterpretation can lead to significant deviations from a valid proof path.
 203 More specifically, the whole-proof generation model may have incorrect beliefs on intermediate
 204 tactic states when generating long proofs. Online reinforcement learning has been proven to be an
 205 effective method for mitigating compounding errors in the extrapolation setting of model inference
 206 (Fujimoto et al., 2019). By continuously interacting with the environment and receiving feedback in
 207 real-time, the model is able to refine its decision-making policy and reduce the bias inducted from
 208 the supervised dataset.

209 4 EXPLORATION-ORIENTED MONTE-CARLO TREE SEARCH

210 4.1 TACTIC-LEVEL TREE ABSTRACTION

211
 212 To implement the tree search method in the whole-proof generation setting, we introduce a proof tree
 213 abstraction to define the tailored state and action space, leveraging a truncate-and-resume mecha-
 214 nism. Roughly following the paradigm of Yao et al. (2023), we begin by decomposing an incomplete
 215 proof into a sequence of tree nodes that correspond to individual proof steps, and then we utilize the



239
240
241
242
243
244
245
246
247
248
249

Figure 2: **Truncate-and-Resume Mechanism in the Expansion Step of MCTS.** (a) After selecting a node, we trace its corresponding incomplete proof code prefix, which includes the file header, initial statement, and successfully applied tactics from the ancestor nodes. (b) The language model then generates the subsequent proof based on this prefix along with a comment block containing the current tactic state. (c) The combined proof code (prefix and newly generated code) is verified by the Lean 4 prover. If no errors are found, the tree-search procedure terminates. If errors are detected, we truncate the newly generated code at the first error message, discard the subsequent code, and parse the successful portion into tactics. (d) Each tactic is added as a new node in the search tree, extending a chain of descendants beneath the selected node. (e) Once the tree updates are complete, the next iteration of expansion begins by selecting an alternative candidate node, which is not limited to leaf nodes. This process repeats until a correct proof is found or the sample budget is exhausted.

250
251
252

partial content stored in these tree nodes to continue the proof generation process. Figure 2 illustrates the process of constructing a proof search tree from whole-proof generation.

253
254
255
256
257
258
259
260
261
262

Truncate: Proof Decomposition into Tree Nodes. We construct the proof search tree at the tactic level, where each tree edge represents a single transition step of the tactic state. Initially, we submit the entire proof the model generated to the Lean prover to parse it into tactics. We then truncate the proof at the earliest verification error, ensuring that all subsequent tactic codes can be successfully applied to advance the proof towards the desired theorem. The tactic codes are segmented into several code fractions, each containing a valid tactic code and its associated chain-of-thought comments, corresponding to a single tree edge that represents a tactic state transition. Through this abstraction, each tactic code is converted into a series of tree nodes, forming a path from the root to a specific node.

263
264
265
266
267
268
269

Resume: Proof Generation from a Tree Node. In Lean 4, different tactics can lead to the same tactic state, meaning each node in our proof tree can correspond to various tactic codes that achieve the same outcome. To handle this, we store a set of these equivalent tactic codes at each node. When the tree search agent expands a node, it randomly selects one tactic to use as a prompt for the language model. This prompt includes the incomplete proof code ending with the chosen tactic and the tactic state information from the Lean prover as a comment block. The fine-tuned model (see Section 3.1) has been trained to recognize and utilize this format, using the incomplete code augmented with tactic state comments to guide subsequent proof generation.

4.2 INTERACTIVE THEOREM PROVING VIA MONTE-CARLO TREE SEARCH

Our proof search tree is developed using the standard Monte-Carlo Tree Search (MCTS) paradigm (MCTS; Coulom, 2006; Browne et al., 2012), which iteratively applies four steps: *Selection*, *Expansion*, *Simulation*, and *Backpropagation*. We integrate the *Simulation* step into *Expansion* because our whole-proof generation model inherently performs a rollout from the expanded node. The detailed design of the algorithm workflow is as follows.

Selection. The selection step, *a.k.a.* the tree policy, starts from the root node and traverses downward to identify a promising node for expansion. The objective of this algorithmic step is to trade off between exploration and exploitation (Kocsis & Szepesvári, 2006). The tree policy at a tree node s is computed by selecting the action that maximizes the value from the set of valid operations:

$$\text{TreePolicy}(s) = \arg \max_{a \in \text{Children}(s) \cup \{\emptyset\}} Q_{\text{UCB}}(s, a), \quad (1)$$

where the action a can be either moving to a child node, denoted by $a \in \text{Children}(s)$, or expanding the current node s , denoted by a special token $a = \emptyset$. This approach uses a technique called *virtual node* (Wang et al., 2023b), which assigns each node an imaginary child to represent the selection of the current node s for expansion. It enables the tree search agent to continually expand non-leaf nodes, as the action space is supported by a generative model whose output scope cannot be determined by a fixed number of trails. The value estimation $Q_{\text{UCB}}(s, a)$ of performing action a on node s is composed by two components:

$$\forall a \in \text{Children}(s) \cup \{\emptyset\}, \quad Q_{\text{UCB}}(s, a) = \underbrace{Q(s, a)}_{\text{Exploitation}} + \underbrace{\text{UCB}(s, a)}_{\text{Exploration}}, \quad (2)$$

where $Q(s, a)$ denotes a sample-based estimation of action values derived from the selection history, functioning as the exploitation component that retrieves high-value candidates from previous trials. $\text{UCB}(s, a)$ denotes the exploration bonus computed by upper confidence bounds (UCB; Auer, 2002), which diminishes with the repeated execution of the state-action pair (s, a) . More specifically, $Q_{\text{UCB}}(s, a)$ stands for an optimistic estimation of $Q(s, a)$ and can serve as an upper bound with high probability. We defer the discussion of detailed settings of node values and UCB bonus to Section 4.3.

Expansion. The next step is invoking the proof generation model to expand the node nominated by the selection phase. Resuming the incomplete proof codes stored on the node designated for expansion, we perform whole-proof generation to propose a series of subsequent tactics and submit the generated proof to Lean prover for verification. Such a trial of proof completion is equivalent to conducting a single rollout of simulation within the standard MCTS framework. When the verification result indicates the proof is complete, the search procedure is ready to be terminated, having found a new proof of the desired theorem. Otherwise, we parse the verification feedback and truncate the generated proof to the assertion of the earliest verification error. The remaining tactics are transformed into a path of nodes to be merged into the search tree (see Figure 2). It is important to note that, because we use the whole-proof generation setting—where the output is an entire proof consisting of a sequence of tactics, rather than just the next tactic—our expansion procedure may insert a path of tree nodes into the search tree during each iteration. This differs from the conventional MCTS designed for competitive games, which typically expands only one layer of children nodes per iteration (Silver et al., 2016; 2018; Schrittwieser et al., 2020).

Backpropagation. The final phase of each tree search iteration is to update value statistics along the selection trajectory from the root to the expanded node, *i.e.*, updating the values associated with the tree policy stated in Eq. (1). Let $\tau = \{\text{root}, s^{(1)}, (s^{(1)}, s^{(2)}), (s^{(2)}, s^{(3)}), \dots, (s^{(|\tau|-1)}, s_t, \emptyset)\}$ denote the selection trajectory of t -th iteration that ends with s_t as the expanding node. We update $Q_{\text{UCB}}(s, a)$ for all $(s, a) \in \tau$ by taking the most recent trajectory reward $R(\tau)$ into account (details refer to Eq. (7)). The extrinsic source of rewards comes from the compiler feedback, specifically assigning a reward of $R_{\text{extrinsic}}(\tau) = 1$ for completed proofs and $R_{\text{extrinsic}}(\tau) = 0$ for unsolved ones. In Section 4.3, we will introduce an intrinsic reward mechanism to augment the reward assignment that enhances the agent’s incentive for exploration.

4.3 INTRINSIC REWARDS FOR MONTE-CARLO TREE SEARCH

In the search problem of formal theorem proving, the extrinsic rewards are extremely sparse, *i.e.*, the search agent only obtains non-zero rewards when the proof is completely solved. More specifically, the proof search process forms a tree structure with only a narrow set of leaves delivering non-zero rewards, which matches a famous hard-exploration case (Krishnamurthy et al., 2016) in the literature of statistical reinforcement learning. To promote exploration in sparse-reward sequential decision making, one classical paradigm is constructing intrinsic rewards (Schmidhuber, 2010) that encourage the agent to not only optimize extrinsic rewards but also acquire general information about the interactive environment (Bellemare et al., 2016; Houthoofd et al., 2016; Pathak et al., 2017; Burda et al., 2019). In this section, we present our intrinsic-reward-driven exploration algorithm, *RMax applied to Tree Search* (RMaxTS), to incorporate reward-free exploration in the proof search problem.

RMax applied to MCTS. We adopt RMax (Brafman & Tenenbholz, 2002), a classical exploration mechanism, to construct intrinsic rewards for Monte-Carlo tree search. The core idea of RMax is to explore a broad coverage of the state space. The agent awards itself a maximal amount of reward upon reaching an unseen state. In the context of proof search, where no extrinsic rewards are provided until the proof is completed, our algorithmic procedure resembles ZeroRMax (Jin et al., 2020), in which the agent’s exploration is driven solely by intrinsic rewards, *i.e.*, setting $R(\tau) = R_{\text{intrinsic}}(\tau)$. The intrinsic reward of a tree expansion step is determined by whether a new node is added to the search tree,

$$R_{\text{intrinsic}}(\tau) = \mathbb{I}[\text{at least one new node is added to the search tree}], \quad (3)$$

where τ denotes the most recent selection trajectory that requires a reward assignment for back-propagation. This exploration strategy prioritizes the expansion of nodes where the prover model generates tactics that lead to a diverse range of tactic states. As multiple Lean codes can result in the same transition of intermediate states, this heuristics can potentially reduce redundant generation and improve sample efficiency.

UCB for Non-stationary Rewards. The common setting of UCB exploration bonus for Monte-Carlo tree search is using UCB1 (Auer et al., 2002):

$$Q_{\text{UCB1}}(s, a) = \frac{W(s, a)}{N(s, a)} + \sqrt{\frac{2 \ln \sum_{a'} N(s, a')}{N(s, a)}}, \quad (4)$$

$$W(s, a) = \sum_{\tau \in \Gamma(s, a)} R(\tau), \quad (5)$$

$$N(s, a) = |\Gamma(s, a)|, \quad (6)$$


where $\Gamma(s, a) = \{\tau \mid (s, a) \in \tau\}$ denotes the list of tree-policy trajectory τ containing (s, a) as an intermediate selection step. To facilitate discussions, we organize the list $\Gamma(s, a) = \{\tau_1, \tau_2, \dots\}$ such that newly collected trajectories have larger subscript indices. In this work, we propose to use an alternative variant of UCB method. Note that the derived intrinsic reward in Eq. (3) is a non-stationary reward signal whose expected value decays with the progress of exploration. That is because it becomes definitely harder to discover new nodes with unseen tactic states as the search tree expands through sophisticated exploration. To tackle the non-stationarity, we consider *discounted upper confidence bounds* (DUCB; Garivier & Moulines, 2011), which uses a discount factor $\gamma \in (0, 1)$ to smoothly drop those outdated feedback records:

$$Q_{\text{DUCB}}(s, a) = \frac{W_\gamma(s, a)}{N_\gamma(s, a)} + \sqrt{\frac{2 \ln \sum_{a'} N_\gamma(s, a')}{N_\gamma(s, a)}}, \quad (7)$$

$$W_\gamma(s, a) = \sum_{t=1}^{N(s, a)} \gamma^{N(s, a)-t} R(\tau_t), \quad (8)$$

$$N_\gamma(s, a) = \sum_{t=0}^{N(s, a)-1} \gamma^t, \quad (9)$$

where newly received feedback would be assigned a larger weight in the value estimation. In practice, we set $\gamma = 0.99$. Note that the role of discount factor γ in DUCB differs from its role in value iteration for infinite-horizon MDPs. The discounting is applied to tree search iterations rather than to the action-step horizon within a single trajectory.



Model	Pass@128	
	miniF2F-test	ProofNet-test
Base (3-shot)	29.7% ± 0.5%	9.7% ± 0.7%
SFT (non-CoT)	49.8% ± 0.3%	15.9% ± 0.5%
SFT (CoT)	50.4% ± 0.4%	15.9% ± 0.6%
RL (non-CoT)	50.5% ± 0.6%	17.5% ± 0.5%
RL (CoT)	51.6% ± 0.5%	18.2% ± 0.5%

Figure 3: **Comparison of model capabilities at different training stages.** "CoT" and "non-CoT" refer to evaluations using two guiding prompts. The shaded region represents the range of standard deviations around the mean values. The notation $\mu \pm \sigma$ indicates the average accuracy μ and the standard deviation σ .

5 EXPERIMENTAL RESULTS

In this section, we evaluate the theorem-proving capabilities of DS-Prover-V1.5 using two distinct benchmarks: miniF2F (Zheng et al., 2022), which encompasses high-school level exercises and competition problems, and ProofNet (Azerbaiyev et al., 2023), which pertains to undergraduate-level theorems. We present the results for both whole-proof generation and Monte-Carlo tree search methodologies. Detailed experiment settings are described in Appendix B.1.

5.1 MAIN RESULTS

Results on miniF2F and ProofNet. Table 1 and 2 provides a comparative analysis of various theorem-proving methods on the miniF2F and ProofNet benchmarks. In the single-pass whole-proof generation setting, DS-Prover-V1.5-RL achieved the highest pass rate at 60.2% on miniF2F-test and at 23.7% on ProofNet-test, significantly outperforming all advanced baselines. When combining DS-Prover-V1.5-RL with RMaxTS, the new state-of-the-art are achieved, solving 62.7% problems from miniF2F-test and 25.3% problems from ProofNet-test.

General Enhancement of Reinforcement Learning. To support the claim that online reinforcement learning from verification feedback generally enhances the model capabilities, we compare our final model to the SFT-only version using a large sample budget. The comparison results are presented as two columns in Table 4. DS-Prover-V1.5-RL consistently outperforms the SFT model across all generation settings, regardless of whether the chain-of-thought strategy is applied. The results also indicate that the improvements gained from conducting online RL is orthogonal to those achieved through RMaxTS, which can be further combined to boost the performance. By integrating both CoT prompting and RMaxTS, DS-Prover-V1.5-RL achieves a pass rate of 62.7% on miniF2F-test. This performance shows a notable 3.7% improvement over the SFT model, highlighting the critical role of reinforcement learning in enhancing the overall effectiveness of the proof completion model.

CoT, non-CoT, and Mixture Strategy. We compare the performance of two generation modes, *i.e.*, non-CoT and CoT, on miniF2F-test dataset. The results, shown in Table 4, indicate that the advantage of CoT over the non-CoT mode is amplified as the sample budget increases. This suggests that the incorporation of natural language chain-of-thought can diversify the planning pathways of theorem proving, potentially leading to a broader range of reasoning strategies and more innovative solutions. Results also show that these two modes have complementary advantages across different problems. The model’s theorem proving strategy in the CoT mode is more systematic and proactive in mathematical thinking, while in the non-CoT mode, the model can efficiently use Lean high-level tactics to solve computational problems that can be addressed within Lean’s automation mechanisms. To leverage these advantages, we consider a mixture strategy, denoted by non-CoT & CoT in Table 4, allocates half of sample budget to the CoT mode and the remains to the non-CoT mode. This simple combination of two guiding prompts shows great promise in further bootstrapping the performance of our proof completion model, achieving a pass rate of 63.5% on miniF2F-test. In Appendix D, we present example problems that illustrate the different advantages of the two generation modes.


Method	Sample budget	miniF2F-test
<i>Single-pass Whole-Proof Generation Methods</i>		
DeepSeek-Prover-V1 (Xin et al., 2024)	128	46.1% \pm 0.5%
	16 \times 4096	50.0%
DS-Prover-V1.5-SFT	128	50.4% \pm 0.4%
	3200	53.3% \pm 0.5%
	16 \times 6400	57.4%
DS-Prover-V1.5-RL	128	51.6% \pm 0.5%
	3200	54.9% \pm 0.7%
	16 \times 6400	60.2%
<i>Tree Search Methods</i>		
GPT-f (Polu et al., 2022)	64 \times 8 \times 512	36.6%
Hypertree Proof Search (Lample et al., 2022)	64 \times 5000	41.0%
Lean-STaR (Lin et al., 2024)	64 \times 1 \times 50	46.3%
InternLM2-Math-Plus-7B (Ying et al., 2024b)	1 \times 32 \times 100	43.4%
InternLM2-StepProver (Wu et al., 2024)	1 \times 32 \times 100	48.8%
	64 \times 32 \times 100	54.5%
DS-Prover-V1.5-SFT + RMaxTS	1 \times 3200	53.5% \pm 0.4%
	16 \times 6400	59.0%
	32 \times 6400 [†]	60.2%
DS-Prover-V1.5-RL + RMaxTS	1 \times 3200	55.0% \pm 0.7%
	16 \times 6400	62.7%
	32 \times 6400 [†]	63.5%

Table 1: Comparison with state-of-the-art methods on the miniF2F-test dataset. Unless otherwise specified, DS-Prover-V1.5-SFT and RL employ CoT mode prompting. The symbol [†] indicates performance using a mixture strategy with two guiding prompts (see Section 5.1 for details). More baseline results are presented in Table 3 in Appendix.

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

Method	Sample budget	ProofNet		
		valid [‡]	test	all
<i>Single-pass Whole-Proof Generation Methods</i>				
DS-Prover-V1.5-SFT	128	19.9% \pm 0.4%	15.9% \pm 0.6%	17.9% \pm 0.3%
	3200	20.7% \pm 0.7%	21.0% \pm 0.9%	20.9% \pm 0.6%
	4 \times 6400	22.2%	23.7%	22.9%
DS-Prover-V1.5-RL	128	20.1% \pm 0.5%	18.2% \pm 0.5%	19.1% \pm 0.4%
	3200	21.4% \pm 0.3%	22.0% \pm 0.5%	21.7% \pm 0.4%
	4 \times 6400	21.6%	23.7%	22.6%
<i>Tree Search Methods</i>				
ReProver (Yang et al., 2023)	-	-	-	13.8%
InternLM2-StepProver (Wu et al., 2024)	1 \times 32 \times 100	-	-	18.1%
DS-Prover-V1.5-SFT + RMaxTS	1 \times 3200	22.2% \pm 0.7%	21.6% \pm 0.2%	21.9% \pm 0.4%
	4 \times 6400	23.8%	25.8%	24.8%
DS-Prover-V1.5-RL + RMaxTS	1 \times 3200	22.0% \pm 0.3%	21.5% \pm 0.8%	21.8% \pm 0.4%
	4 \times 6400	25.4%	25.3%	25.3%

Table 2: Comparing with state-of-the-arts on the ProofNet dataset. [‡] Note that the validation set of ProofNet is used to perform expert iteration in supervised fine-tuning.

486
487
488
489
490
491
492
493
494
495
496

	Sample budget	miniF2F-test
Single-Pass Generation	4×6400	$58.4\% \pm 0.5\%$
	16×6400	60.2%
UCT (without $R_{\text{intrinsic}}$)	4×6400	$58.2\% \pm 0.3\%$
	16×6400	61.1%
RMaxTS (DUCB \rightarrow UCB1)	4×6400	$58.6\% \pm 0.3\%$
	16×6400	60.7%
RMaxTS (without tactic state)	4×6400	$58.4\% \pm 0.3\%$
	16×6400	61.1%
RMaxTS	4×6400	$59.6\% \pm 0.6\%$
	16×6400	62.7%

Figure 4: A modular ablation study examining the algorithmic design of RMaxTS. The experiments are conducted on the miniF2F-test dataset with DS-Prover-V1.5-RL using the CoT mode. The left panel presents the curves of Pass@K accuracy within 6400 generation samples. The results with a larger sample size are presented in the right panel.

501
502
503

5.2 ABLATION STUDIES ON RMAXTS

504
505
506
507
508
509
510
511
512
513
514
515
516
517

Intrinsic Rewards and Discounted UCB. We investigate the effectiveness of two core components of RMaxTS, *i.e.*, the intrinsic rewards defined in Eq. (3) and the discounted upper confidence bound stated in Eq. (7). We start with a baseline implementing the standard UCT algorithm (Kocsis & Szepesvári, 2006) without intrinsic rewards, in which the exploration is driven exclusively by the UCB bonus. Note that, since no non-zero rewards are provided for this baseline, all variants of the UCB formula become equivalent, as node selection is determined solely by visitation counts. The experimental results in Figure 4 show that, in the absence of intrinsic rewards, the performance of UCT (without $R_{\text{intrinsic}}$) degenerates into a level comparable to that of non-search methods. Furthermore, we consider RMaxTS using the standard UCB1 (refer to Eq. (4)) instead of the discounted UCB, denoted by RMaxTS (DUCB \rightarrow UCB1). The results indicate that the performance of RMaxTS with UCB1 bonus is also moderate, comparable to that of UCT (without $R_{\text{intrinsic}}$). That is because UCB1 is designed to guarantee asymptotic performance through exhausted exploration (Auer et al., 2002) assuming the sample size to be sufficiently large. In contrast, the discounted UCB can accelerate the value propagation of non-stationary intrinsic rewards.

518
519
520
521
522
523
524
525
526

Guidance of Tactic State Information. When expanding a tree node, we concatenate the intermediate tactic state information as a comment block to the incomplete code to guide the proof completion. With the provided auxiliary information, the proof completion model can enhance its internal representation of the tactic state, offering intermediate guidance for long-horizon planning. To demonstrate this advantage, we present experiments on RMaxTS that performs code completion directly from the raw incomplete code without accessing tactic state information, denoted by RMaxTS (without tactic state) in Figure 4. The results indicate that the performance gain from applying tree search becomes moderate in the absence of tactic state information, especially when tackling hard problems that require a large amount of samples.

527
528
529

6 CONCLUSION

530
531
532
533
534
535
536
537
538
539

The framework of DS-Prover-V1.5 is designed to establish an AlphaZero-like pipeline for formal theorem proving. The use of expert iteration and synthetic data mirrors the core trial-and-error loop of reinforcement learning, with the compiler oracle serving as the world model to provide environmental supervision. Within the RL paradigm, the integrated tree search module has proven to be highly effective in advancing superhuman performance across various domains (Silver et al., 2016; Fawzi et al., 2022; Lutz et al., 2023). A promising future direction is training a critic model to assess incomplete proofs and prune search branches of proofs. Such a partial-proof critic model would implicitly perform temporal credit assignment (Sutton, 1984), decomposing proof-level feedback into step-wise value differences (Arjona-Medina et al., 2019). Developing critic models for assessing long planning paths and providing guidance rewards presents a crucial and challenging problem (Ng & Russell, 2000; Sorg et al., 2010) that warrants further investigation.

REFERENCES

- 540
541
542 Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob
543 McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. In
544 *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pp.
545 5055–5065, 2017.
- 546 Jose A Arjona-Medina, Michael Gillhofer, Michael Widrich, Thomas Unterthiner, Johannes Brand-
547 stetter, and Sepp Hochreiter. Rudder: Return decomposition for delayed rewards. In *Proceed-
548 ings of the 33rd International Conference on Neural Information Processing Systems*, pp. 13566–
549 13577, 2019.
- 550 Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine
551 Learning Research*, 3(Nov):397–422, 2002.
- 552 Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit
553 problem. *Machine learning*, 47:235–256, 2002.
- 554 Eser Aygün, Ankit Anand, Laurent Orseau, Xavier Glorot, Stephen M McAleer, Vlad Firoiu, Lei M
555 Zhang, Doina Precup, and Shihab Mourad. Proving theorems using incremental learning and
556 hindsight experience replay. In *International Conference on Machine Learning*, pp. 1198–1210.
557 PMLR, 2022.
- 558 Zhangir Azerbayev, Bartosz Piotrowski, Hailey Schoelkopf, Edward W Ayers, Dragomir Radev, and
559 Jeremy Avigad. Proofnet: Autoformalizing and formally proving undergraduate-level mathemat-
560 ics. *arXiv preprint arXiv:2302.12433*, 2023.
- 561 Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen Marcus McAleer,
562 Albert Q Jiang, Jia Deng, Stella Biderman, and Sean Welleck. Llemma: An open language model
563 for mathematics. In *The Twelfth International Conference on Learning Representations*, 2024.
- 564 Marc G Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Rémi
565 Munos. Unifying count-based exploration and intrinsic motivation. In *Proceedings of the 30th
566 International Conference on Neural Information Processing Systems*, pp. 1479–1487, 2016.
- 567 Ronen I Brafman and Moshe Tennenholtz. R-max-a general polynomial time algorithm for near-
568 optimal reinforcement learning. *Journal of Machine Learning Research*, 3(Oct):213–231, 2002.
- 569 Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp
570 Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of
571 monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in
572 games*, 4(1):1–43, 2012.
- 573 Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network
574 distillation. In *Seventh International Conference on Learning Representations*, pp. 1–17, 2019.
- 575 Guillaume MJ B Chaslot, Mark HM Winands, and H Jaap van Den Herik. Parallel monte-carlo
576 tree search. In *Computers and Games: 6th International Conference, CG 2008, Beijing, China,
577 September 29-October 1, 2008. Proceedings 6*, pp. 60–71. Springer, 2008.
- 578 Rémi Coulom. Efficient selectivity and backup operators in Monte-Carlo tree search. In *Interna-
579 tional conference on computers and games*, pp. 72–83. Springer, 2006.
- 580 Maxwell Crouse, Ibrahim Abdelaziz, Bassem Makni, Spencer Whitehead, Cristina Cornelio, Pavan
581 Kapanipathi, Kavitha Srinivas, Veronika Thost, Michael Witbrock, and Achille Fokoue. A deep
582 reinforcement learning approach to first-order logic theorem proving. In *Proceedings of the AAAI
583 Conference on Artificial Intelligence*, volume 35, pp. 6279–6287, 2021.
- 584
585
586
587
588
589
590
591
592
593 Alhussein Fawzi, Mateusz Malinowski, Hamza Fawzi, and Omar Fawzi. Learning dynamic poly-
nomial proofs. In *Proceedings of the 33rd International Conference on Neural Information Pro-
cessing Systems*, pp. 4179–4188, 2019.

- 594 Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Moham-
595 madamin Barekatin, Alexander Novikov, Francisco J R Ruiz, Julian Schrittwieser, Grzegorz
596 Swirszcz, et al. Discovering faster matrix multiplication algorithms with reinforcement learning.
597 *Nature*, 610(7930):47–53, 2022.
- 598
599 Guhao Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. Towards revealing
600 the mystery behind chain of thought: A theoretical perspective. In *Thirty-seventh Conference on*
601 *Neural Information Processing Systems*, volume 36, 2023.
- 602 Scott Fujimoto, David Meger, and Doina Precup. Off-policy deep reinforcement learning without
603 exploration. In *International conference on machine learning*, pp. 2052–2062. PMLR, 2019.
- 604
605 Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for switching bandit
606 problems. In *International conference on algorithmic learning theory*, pp. 174–188. Springer,
607 2011.
- 608 Thibault Gauthier. Deep reinforcement learning for synthesizing functions in higher-order logic.
609 *EPiC Series in Computing*, 73:230–248, 2020.
- 610
611 Rein Houthoofd, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. Vime:
612 variational information maximizing exploration. In *Proceedings of the 30th International Con-*
613 *ference on Neural Information Processing Systems*, pp. 1117–1125, 2016.
- 614 Jiewen Hu, Thomas Zhu, and Sean Welleck. minictx: Neural theorem proving with (long-)contexts,
615 2024.
- 616
617 Albert Q Jiang, Wenda Li, Szymon Tworkowski, Konrad Czechowski, Tomasz Odrzygóźdź, Piotr
618 Miłoś, Yuhuai Wu, and Mateja Jamnik. Thor: wielding hammers to integrate language models
619 and automated theorem provers. In *Proceedings of the 36th International Conference on Neural*
620 *Information Processing Systems*, pp. 8360–8373, 2022a.
- 621
622 Albert Qiaochu Jiang, Sean Welleck, Jin Peng Zhou, Timothee Lacroix, Jiacheng Liu, Wenda Li,
623 Mateja Jamnik, Guillaume Lample, and Yuhuai Wu. Draft, sketch, and prove: Guiding formal
624 theorem provers with informal proofs. In *The Eleventh International Conference on Learning*
625 *Representations*, 2022b.
- 626
627 Chi Jin, Akshay Krishnamurthy, Max Simchowitz, and Tiancheng Yu. Reward-free exploration
628 for reinforcement learning. In *International Conference on Machine Learning*, pp. 4870–4879.
PMLR, 2020.
- 629
630 Levente Kocsis and Csaba Szepesvári. Bandit based Monte-Carlo planning. In *European conference*
631 *on machine learning*, pp. 282–293. Springer, 2006.
- 632
633 Akshay Krishnamurthy, Alekh Agarwal, and John Langford. PAC reinforcement learning with rich
634 observations. In *Proceedings of the 30th International Conference on Neural Information Pro-*
cessing Systems, pp. 1848–1856, 2016.
- 635
636 Mitsuru Kusumoto, Keisuke Yahata, and Masahiro Sakai. Automated theorem proving in intuition-
637 istic propositional logic by deep reinforcement learning. *arXiv preprint arXiv:1811.00796*, 2018.
- 638
639 Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E.
640 Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model
641 serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating*
Systems Principles, 2023.
- 642
643 Guillaume Lample, Marie-Anne Lachaux, Thibaut Lavril, Xavier Martinet, Amaury Hayat, Gabriel
644 Ebner, Aurélien Rodriguez, and Timothée Lacroix. Hypertree proof search for neural theorem
645 proving. In *Proceedings of the 36th International Conference on Neural Information Processing*
646 *Systems*, pp. 26337–26349, 2022.
- 647
647 Leanprover Community. A read-eval-print-loop for Lean 4. <https://github.com/leanprover-community/repl>, 2023.

- 648 Jannis Limperg and Asta Halkjær From. Aesop: White-box best-first proof search for lean. In
649 *Proceedings of the 12th ACM SIGPLAN International Conference on Certified Programs and*
650 *Proofs*, pp. 253–266, 2023.
- 651 Haohan Lin, Zhiqing Sun, Yiming Yang, and Sean Welleck. Lean-star: Learning to interleave
652 thinking and proving. *arXiv preprint arXiv:2407.10040*, 2024.
- 653 Isaac D Lutz, Shunzhi Wang, Christoffer Norn, Alexis Courbet, Andrew J Borst, Yan Ting Zhao,
654 Annie Dosey, Longxing Cao, Jinwei Xu, Elizabeth M Leaf, et al. Top-down design of protein
655 architectures with reinforcement learning. *Science*, 380(6642):266–273, 2023.
- 656 Mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN*
657 *International Conference on Certified Programs and Proofs*, pp. 367–381. Association for Com-
658 puting Machinery, 2020.
- 659 Jack McKeown and Geoff Sutcliffe. Reinforcement learning for guiding the e theorem prover. In
660 *The International FLAIRS Conference Proceedings*, volume 36, 2023.
- 661 Leonardo de Moura and Sebastian Ullrich. The lean 4 theorem prover and programming language.
662 In *Automated Deduction—CADE 28: 28th International Conference on Automated Deduction,*
663 *Virtual Event, July 12–15, 2021, Proceedings 28*, pp. 625–635. Springer, 2021.
- 664 Andrew Y Ng and Stuart J Russell. Algorithms for inverse reinforcement learning. In *Proceedings*
665 *of the Seventeenth International Conference on Machine Learning*, pp. 663–670, 2000.
- 666 OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- 667 Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration
668 by self-supervised prediction. In *International conference on machine learning*, pp. 2778–2787.
669 PMLR, 2017.
- 670 Lawrence C. Paulson. *Isabelle a Generic Theorem Prover*. Springer Verlag, 1994.
- 671 Stanislas Polu and Ilya Sutskever. Generative language modeling for automated theorem proving.
672 *arXiv preprint arXiv:2009.03393*, 2020.
- 673 Stanislas Polu, Jesse Michael Han, Kunhao Zheng, Mantas Baksys, Igor Babuschkin, and Ilya
674 Sutskever. Formal mathematics statement curriculum learning. *arXiv preprint arXiv:2202.01344*,
675 2022.
- 676 Michael Rawson and Giles Reger. A neurally-guided, parallel theorem prover. In *Frontiers of*
677 *Combining Systems: 12th International Symposium, FroCoS 2019, London, UK, September 4-6,*
678 *2019, Proceedings 12*, pp. 40–56. Springer, 2019.
- 679 Michael Rawson and Giles Reger. lazycop: Lazy paramodulation meets neurally guided search. In
680 *Automated Reasoning with Analytic Tableaux and Related Methods: 30th International Confer-*
681 *ence, TABLEUX 2021, Birmingham, UK, September 6–9, 2021, Proceedings 30*, pp. 187–199.
682 Springer, 2021.
- 683 Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and struc-
684 tured prediction to no-regret online learning. In *Proceedings of the fourteenth international con-*
685 *ference on artificial intelligence and statistics*, pp. 627–635. JMLR Workshop and Conference
686 Proceedings, 2011.
- 687 Jürgen Schmidhuber. Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE*
688 *transactions on autonomous mental development*, 2(3):230–247, 2010.
- 689 Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon
690 Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari,
691 go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- 692 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
693 optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 694
695
696
697
698
699
700
701

- 702 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Mingchuan Zhang, YK Li, Y Wu,
703 and Daya Guo. DeepSeekMath: Pushing the limits of mathematical reasoning in open language
704 models. *arXiv preprint arXiv:2402.03300*, 2024.
- 705
- 706 David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche,
707 Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering
708 the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- 709
- 710 David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez,
711 Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, et al. A general reinforcement
712 learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–
713 1144, 2018.
- 714 Jonathan Sorg, Satinder Singh, and Richard L Lewis. Reward design via online gradient ascent.
715 In *Proceedings of the 23rd International Conference on Neural Information Processing Systems-*
716 *Volume 2*, pp. 2190–2198, 2010.
- 717
- 718 Richard Stuart Sutton. Temporal credit assignment in reinforcement learning. *Phd thesis, University*
719 *of Massachusetts*, 1984.
- 720
- 721 Amitayush Thakur, Yeming Wen, and Swarat Chaudhuri. A language-agent approach to formal
722 theorem-proving. In *The 3rd Workshop on Mathematical Reasoning and AI at NeurIPS*, 2023.
- 723
- 724 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
725 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural informa-*
726 *tion processing systems*, 30, 2017.
- 727
- 728 Haiming Wang, Huajian Xin, Chuanyang Zheng, Zhengying Liu, Qingxing Cao, Yinya Huang, Jing
729 Xiong, Han Shi, Enze Xie, Jian Yin, et al. Lego-prover: Neural theorem proving with growing
730 libraries. In *The Twelfth International Conference on Learning Representations*, 2023a.
- 731
- 732 Haiming Wang, Ye Yuan, Zhengying Liu, Jianhao Shen, Yichun Yin, Jing Xiong, Enze Xie, Han
733 Shi, Yujun Li, Lin Li, et al. DT-solver: Automated theorem proving with dynamic-tree sam-
734 pling guided by proof-level value function. In *Proceedings of the 61st Annual Meeting of the*
735 *Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12632–12646, 2023b.
- 736
- 737 Ruida Wang, Jipeng Zhang, Yizhen Jia, Rui Pan, Shizhe Diao, Renjie Pi, and Tong Zhang. Theorem-
738 llama: Transforming general-purpose llms into lean4 experts. *arXiv preprint arXiv:2407.03203*,
739 2024.
- 740
- 741 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H Chi,
742 Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language
743 models. In *Proceedings of the 36th International Conference on Neural Information Processing*
744 *Systems*, pp. 24824–24837, 2022.
- 745
- 746 Sean Welleck and Rahul Saha. llmstep: Llm proofstep suggestions in lean. In *The 3rd Workshop on*
747 *Mathematical Reasoning and AI at NeurIPS’23*, 2023.
- 748
- 749 Zijian Wu, Jiayu Wang, Dahua Lin, and Kai Chen. Lean-github: Compiling github lean repositories
750 for a versatile lean prover. *arXiv preprint arXiv:2407.17227*, 2024.
- 751
- 752 Huajian Xin, Daya Guo, Zhihong Shao, Zhizhou Ren, Qihao Zhu, Bo Liu, Chong Ruan, Wenda Li,
753 and Xiaodan Liang. Deepseek-prover: Advancing theorem proving in llms through large-scale
754 synthetic data. *arXiv preprint arXiv:2405.14333*, 2024.
- 755
- 756 Kaiyu Yang. minif2f-lean4. <https://github.com/yangky11/minif2f-lean4>, 2023.
- 757
- 758 Kaiyu Yang, Aidan M Swope, Alex Gu, Rahul Chalamala, Peiyang Song, Shixing Yu, Saad Godil,
759 Ryan Prenger, and Anima Anandkumar. Leandojo: theorem proving with retrieval-augmented
760 language models. In *Proceedings of the 37th International Conference on Neural Information*
761 *Processing Systems*, pp. 21573–21612, 2023.

- 756 Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik
757 Narasimhan. Tree of thoughts: deliberate problem solving with large language models. In *Pro-
758 ceedings of the 37th International Conference on Neural Information Processing Systems*, pp.
759 11809–11822, 2023.
- 760 Huaiyuan Ying, Zijian Wu, Yihan Geng, Jiayu Wang, Dahua Lin, and Kai Chen. Lean workbook:
761 A large-scale lean problem set formalized from natural language math problems. *arXiv preprint
762 arXiv:2406.03847*, 2024a.
- 763 Huaiyuan Ying, Shuo Zhang, Linyang Li, Zhejian Zhou, Yunfan Shao, Zhaoye Fei, Yichuan Ma,
764 Jiawei Hong, Kuikun Liu, Ziyi Wang, et al. Internlm-math: Open math large language models
765 toward verifiable reasoning. *arXiv preprint arXiv:2402.06332*, 2024b.
- 766 Kaiyan Zhang, Biqing Qi, and Bowen Zhou. Towards building specialized generalist ai with system
767 1 and system 2 fusion. *arXiv preprint arXiv:2407.08642*, 2024.
- 768 Xueliang Zhao, Wenda Li, and Lingpeng Kong. Decomposing the enigma: Subgoal-based demon-
769 stration learning for formal theorem proving. *arXiv preprint arXiv:2305.16366*, 2023.
- 770 Chuanyang Zheng, Haiming Wang, Enze Xie, Zhengying Liu, Jiankai Sun, Huajian Xin, Jianhao
771 Shen, Zhenguo Li, and Yu Li. Lyra: Orchestrating dual correction in automated theorem proving.
772 *arXiv preprint arXiv:2309.15806*, 2023.
- 773 Kunhao Zheng, Jesse Michael Han, and Stanislas Polu. miniF2F: a cross-system benchmark for
774 formal olympiad-level mathematics. In *International Conference on Learning Representations*,
775 2022.
- 776 Qihao Zhu, Daya Guo, Zhihong Shao, Dejian Yang, Peiyi Wang, Runxin Xu, Y Wu, Yukun Li,
777 Huazuo Gao, Shirong Ma, et al. DeepSeek-Coder-V2: Breaking the barrier of closed-source
778 models in code intelligence. *arXiv preprint arXiv:2406.11931*, 2024.
- 779 Zsolt Zombori, Adrián Csiszárík, Henryk Michalewski, Cezary Kaliszyk, and Josef Urban. Towards
780 finding longer proofs. In *Automated Reasoning with Analytic Tableaux and Related Methods:
781 30th International Conference, TABLEUX 2021, Birmingham, UK, September 6–9, 2021, Pro-
782 ceedings 30*, pp. 167–186. Springer, 2021.

787 A IMPLEMENTATION DETAILS

788 A.1 PRE-TRAINING

789 To enhance our language model’s proficiency in generating formal proofs and reasoning through
790 mathematical language, we further pre-train our base model (Shao et al., 2024). This refinement
791 involved training on high-quality datasets that include both code and natural language mathematical
792 content. We specifically focused on formal languages widely used in proof assistants, such as Lean,
793 Isabelle, and Metamath. We designate this improved model as DS-Prover-V1.5-Base.

794 A.2 SUPERVISED FINE-TUNING

795 **Data Curation.** We develop a comprehensive Lean 4 code completion dataset for the supervised
796 fine-tuning. This dataset includes synthetic proof code derived from a wide range of formal theo-
797 rems. These theorems are sourced from various projects, such as the standard Lean 4 math library
798 Mathlib4 (Mathlib Community, 2020), synthetic theorems from DeepSeek-Prover-V1 (Xin et al.,
799 2024) and Lean Workbook (Ying et al., 2024a), and validation sets from the miniF2F (Zheng et al.,
800 2022) and ProofNet (Azerbayev et al., 2023) benchmarks. To augment the formal proof data, we
801 employed an expert iteration process (Polu & Sutskever, 2020). This involves generating proofs
802 using the language model, verifying the generated proof data, retraining the model with the verified
803 data, and then using the optimized model to generate additional proof data. Between each itera-
804 tion, we use DeepSeek-Coder V2 236B (Zhu et al., 2024) to annotate the thought process before
805 the proof code as comments. Finally, we tailor these data for the truncate-and-resume mechanism
806 for Monte-Carlo Tree Search (details in Section 4.1). The resulting proof dataset consists of 9,645k
807 sequences.

Training Setting. We conduct supervised fine-tuning based on the pre-trained model and train for 9B tokens, using a batch size of 2,048 and a constant learning rate of $1e-4$. The training process begins with 100 warm-up steps to stabilize the learning dynamics. Training examples are randomly concatenated to form sequences, with a maximum context length of 4,096 tokens.

A.3 REINFORCEMENT LEARNING FROM PROOF ASSISTANT FEEDBACK

Prompts. In the reinforcement learning stage, we use a subset of theorem statements from the supervised fine-tuning dataset as training prompts. We select theorems for which DS-Prover-V1.5-SFT has a moderate success rate in generating correct proofs upon multiple attempts. This ensures that the model has room for improvement while still being able to receive positive feedback. After filtering, we retain approximately 4.5k unique theorem statements. Each theorem is prefixed with both CoT and non-CoT guiding prompts to enhance the model’s proof generation capabilities in both modes.

Rewards. When training LLMs via RL, a trained reward model typically provides feedback signals. In contrast, formal theorem proving benefits from the rigorous verification of generated proofs by proof assistants, offering a significant advantage. Specifically, each generated proof receives a reward of 1 if verified as correct, and 0 otherwise. While this binary reward signal is accurate, it is also sparse, especially for theorems that are challenging for the supervised fine-tuned model. To mitigate this sparsity, we select training prompts that are challenging yet achievable for the supervised fine-tuned model, as described above.

Reinforcement Learning Algorithm. We employ the Group Relative Policy Optimization (GRPO; Shao et al., 2024) as our RL algorithm, which has demonstrated superior effectiveness and efficiency compared to PPO (Schulman et al., 2017), primarily because it eliminates the necessity of training an additional critic model. Specifically, GRPO samples a group of candidate proofs for each theorem prompt and optimizes the model based on the relative rewards of the outputs within the group. Our prompt selection strategy is designed to likely include both correct and incorrect proofs among the candidates, aligning well with the group-relative nature of GRPO and thereby enhancing the training process.

Training Setting. We conduct RL training based on the SFT model, which serves as both the initial model and the reference model for imposing the Kullback-Leibler (KL) divergence penalty. We use a constant learning rate of $5e-6$, and the KL penalty coefficient is set to 0.02. For each theorem, we sample a group of 32 candidate proofs, with maximum length set to 2,048. The training batch size is configured to 512.

A.4 PARALLELIZATION OF MONTE-CARLO TREE SEARCH

To enhance the efficiency of Monte-Carlo Tree Search (MCTS), we implement several established parallelization techniques as described by Chaslot et al. (2008).

- **Root Parallelization:** We deploy 256 MCTS runners per node, with one language model per GPU and a batch size of 512 for proof generation. The Lean prover is invoked through REPL and executed on a cluster with thousands of CPU cores, where each proof verification task is handled by an individual process, created and terminated in a sandbox. Both proof generation by language models and verification by Lean provers are handled asynchronously. This setup allows MCTS runners to perform concurrent tree search operations, significantly accelerating the process.
- **Tree Parallelization:** We manage each search tree with 32 thread workers to parallelize the tree iteration steps. This method effectively schedules and balances the tasks of proof generation and Lean verification. Each thread worker iteratively performs the tree search loop by selecting a candidate node for expansion, invoking the language model to generate the proof, verifying the generated proof with the Lean prover, and performing backpropagation.
- **Virtual Loss:** To encourage diverse node selection among concurrent thread workers, we assign a virtual reward $R(\tau) = 0$ for ongoing iterations. This involves backpropagating a

reward of 0 temporarily and updating it to the true reward upon completion. This strategy promotes exploration of different nodes for expansion, thereby enhancing the overall search efficiency.

A.5 COMPARISON WITH EXISTING METHODS

In this section, we compare our proposed proof tree search method, which introduces a novel truncate-and-resume mechanism for whole-proof generation, with existing approaches. Current methods for using language models in formal mathematics proof search generally fall into two main strategies:

- **Multi-pass proof-step generation:** This strategy breaks down the proving process into multiple episodes of tactic generation and verification, typically following a **tree search** pattern. It involves generating and verifying one tactic at a time, repeating the process for the next tactic until no proof goals remain. Notable examples include GPT-f (Polu & Sutskever, 2020; Polu et al., 2022), Thor (Jiang et al., 2022a), ReProver (Yang et al., 2023), Hypertree Proof Search (Lample et al., 2022), and InternLM2-StepProver (Wu et al., 2024).
- **Single-pass whole-proof generation:** This approach generates and verify an entire proof in one attempt. If the proof is incorrect, the model generates a new proof in the next attempt. Methods in this category include DSP (Jiang et al., 2022b), Subgoal-Prover Zhao et al. (2023), LEGO-Prover (Wang et al., 2023a), Lyra (Zheng et al., 2023), and miniCTX (Hu et al., 2024).

Our proof tree search method uniquely bridges these two strategies, offering a novel hybrid approach. It starts with whole-proof generation, similar to the single-pass approach, but extends this by implementing a sophisticated truncate-and-resume mechanism. This process involves truncating the generated proof to its successful initial segment, parsing this segment into individual tactics, and resuming the tree search from this point. This iterative process effectively implements a Monte-Carlo Tree Search, seamlessly integrating single-pass whole-proof generation with multi-pass proof-step generation. Consequently, we can train a single model with nearly identical objectives to support both strategies simultaneously. Our experimental results demonstrate that this unified approach achieves superior performance in both settings. By combining the strengths of existing methods and introducing innovative techniques, our method offers a more versatile and effective solution for formal mathematics proof search, potentially paving the way for future advancements in this field.

B EXPERIMENT SETTINGS

B.1 EVALUATION

Benchmarks. We evaluate theorem-proving performance on the following benchmarks to compare model capabilities after each training stage:

- **MiniF2F** (Zheng et al., 2022) focuses on formal problem-solving skills for high-school level exercises and competitions, such as AMC, AIME, and IMO, with an emphasis on algebra and number theory. The benchmark includes 244 validation and 244 test problems, originally in Lean 3 and manually converted to Lean 4.9.0, based on the version provided by Yang (2023).
- **ProofNet** (Azerbaiyev et al., 2023) evaluates formal theorem-proving capabilities at the undergraduate level in mathematics. It comprises 185 validation and 186 test problems from widely-used undergraduate textbooks, covering real and complex analysis, linear algebra, abstract algebra, and topology. These problems were initially in Lean 3 and manually converted to Lean 4.9.0.

Prompting Configurations. For each proof attempt of DS-Prover-V1.5-Base, we independently sample three proof demonstrations from the validation set to construct the few-shot prompts. For the miniF2F benchmark, we use human-written proofs from Yang (2023), while for the ProofNet benchmark, we use correct proofs generated by DS-Prover-V1.5-RL as few-shot demonstrations. For DS-Prover-V1.5-SFT and DS-Prover-V1.5-RL, we employ two types of guiding prompts: one

918 that encourages chain-of-thought (CoT) reasoning before each proof step, and one that does not
 919 (non-CoT). Detailed examples are provided in Appendix C.
 920

921 **Evaluation Settings.** We evaluate theorem-proving performance using the $\text{pass}@K$ accuracy met-
 922 ric, which measures the model’s success in generating a correct proof within K attempts. Each
 923 model is deployed on a single A100-40G GPU, utilizing the vLLM framework (Kwon et al., 2023)
 924 for sample generation. The sampling parameters are set with a temperature of 1, a top-p value of
 925 0.95, and a maximum token limit of 2,048. The generated proofs are then verified using the Lean
 926 4 theorem prover. For this verification, we import Mathlib4 (Mathlib Community, 2020) and Ae-
 927 sop (Limperg & From, 2023) to access predefined premises and tactics. The verification process is
 928 subject to a time limit of 300 seconds.

929 **Baselines.** We present a comparative analysis of DS-Prover-V1.5 against previous state-of-the-art
 930 language models, highlighting its performance and advancements.
 931

- 932 • **General-purpose Models: GPT-3.5** and **GPT-4** (OpenAI, 2023) are advanced generative
 933 AI models developed by OpenAI, known for their effectiveness across diverse tasks, in-
 934 cluding code generation. Despite not being specifically designed for theorem proving, their
 935 extensive parameter scales provide significant capabilities. The evaluation of these models
 936 in formal theorem proving is facilitated by **COPRA** (Thakur et al., 2023), an in-context
 937 learning agent that leverages these large language models to propose tactic applications.
 938 Additionally, we examine **Llemma** (Azerbaiyev et al., 2024), a series of language models
 939 trained on extensive general mathematical corpora, commonly used as the base model for
 940 formal theorem proving.
- 941 • **Specialized Models for Formal Mathematics: GPT-f** (Polu & Sutskever, 2020; Polu
 942 et al., 2022) represents an initial effort to apply Transformers (Vaswani et al., 2017) to
 943 proof-step generation for theorem proving tasks, utilizing a best-first search module to con-
 944 struct complete proofs. Subsequent advancements include **ReProver** (Yang et al., 2023),
 945 **LLMStep** (Welleck & Saha, 2023), and **Lean-STaR** (Lin et al., 2024). **Hypertree Proof**
 946 **Search** (Lample et al., 2022) explores the use of Monte Carlo tree search in formal the-
 947 orem proving using Lean. Concurrent works, **InternLM2-Math** (Ying et al., 2024b) and
 948 **InternLM2-StepProver** (Wu et al., 2024), also demonstrate outstanding performance.

949 **Metric.** We compare the performance of DS-Prover-V1.5 with state-of-the-art models using the
 950 $\text{pass}@K$ accuracy metric, which evaluates the model’s ability to generate a correct proof within
 951 K attempts. We display the sample budget K according to the the following rules to align the
 952 computation budget across different generation schemes.
 953

- 954 • For single-pass sampling methods, we define the sample budget K as the total number
 955 of proofs generated, with large values of K factorized for the ease of comparison to tree
 956 search methods.
- 957 • For best-first-search methods, following the notation of Azerbaiyev et al. (2024), we present
 958 $K = N \times S \times T$ where N denotes the number of best-first-search attempts, S denotes the
 959 number of tactics generated for each expansion, and T denotes the number of expansion
 960 iterations.
- 961 • For tree search methods, *e.g.*, RMaxTS and HTPS (Lample et al., 2022), we present $K =$
 962 $N \times T$ where N denotes the number of tree search attempts, and T denotes the number of
 963 model generations invoked in tree expansions.
 964
 965
 966
 967
 968
 969
 970
 971

B.2 SUPPLEMENTARY EXPERIMENT RESULTS

Method	Sample budget	miniF2F-test	
<i>Single-pass Whole-Proof Generation Methods</i>			
TheoremLlama [59]	128	33.6%	
DeepSeek-Prover-V1 [63]	128	46.1% \pm 0.5%	
	16 \times 4096	50.0%	
DS-Prover-V1.5-Base	128	29.7% \pm 0.5%	
	3200	39.2%	
	6400	42.2%	
DS-Prover-V1.5-SFT	32	48.2% \pm 0.6%	
	64	49.6% \pm 0.7%	
	128	50.4% \pm 0.4%	
	3200	53.3% \pm 0.5%	
	4 \times 6400	55.8% \pm 0.7%	
DS-Prover-V1.5-RL	16 \times 6400	57.4%	
	32	50.0% \pm 0.5%	
	64	50.7% \pm 0.4%	
	128	51.6% \pm 0.5%	
	3200	54.9% \pm 0.7%	
	4 \times 6400	58.4% \pm 0.6%	
	16 \times 6400	60.2%	
	<i>Tree Search Methods</i>		
	COPRA (Code Llama) [55]	1 \times 500	5.7%
	COPRA (GPT-3.5) [55]	1 \times 60	9.0%
COPRA (GPT-4) [55]	1 \times 60	26.6%	
Llemma-7B [7]	1 \times 32 \times 100	26.2%	
Llemma-34B [7]	1 \times 32 \times 100	25.8%	
ReProver [65]	-	26.5%	
LLMStep [61]	1 \times 32 \times 100	27.9%	
GPT-f [43]	64 \times 8 \times 512	36.6%	
Hypertree Proof Search [30]	64 \times 5000	41.0%	
Lean-STaR [33]	64 \times 1 \times 50	46.3%	
InternLM2-Math-7B [68]	1 \times 32 \times 100	30.3%	
InternLM2-Math-Plus-7B [68]	1 \times 32 \times 100	43.4%	
InternLM2-StepProver [62]	1 \times 32 \times 100	48.8%	
	64 \times 32 \times 100	54.5%	
DS-Prover-V1.5-SFT + RMaxTS	1 \times 3200	53.5% \pm 0.4%	
	4 \times 6400	56.3% \pm 0.3%	
	16 \times 6400	59.0%	
	32 \times 6400 [†]	60.2%	
DS-Prover-V1.5-RL + RMaxTS	1 \times 3200	55.0% \pm 0.7%	
	4 \times 6400	59.6% \pm 0.6%	
	16 \times 6400	62.7%	
	32 \times 6400 [†]	63.5%	

Table 3: Comparison with state-of-the-art methods on the miniF2F-test dataset. The notation $\mu \pm \sigma$ denotes the average accuracy μ and the standard deviation σ . Unless otherwise specified, DS-Prover-V1.5-Base results are based on 3-shot prompting, while DS-Prover-V1.5-SFT and RL employ CoT mode prompting. The symbol [†] indicates performance using a mixture strategy with two guiding prompts.

	Prompt mode	Sample budget	DS-Prover-V1.5	
			SFT	RL
Single-Pass Generation	non-CoT	4×6400	$54.7\% \pm 0.4\%$	$56.5\% \pm 0.5\%$
		16×6400	56.1%	57.4%
	CoT	4×6400	$55.8\% \pm 0.7\%$	$58.4\% \pm 0.5\%$
		16×6400	57.4%	60.2%
	non-CoT & CoT	$(2 + 2) \times 6400$	$56.1\% \pm 0.8\%$	$58.3\% \pm 0.6\%$
		$(8 + 8) \times 6400$ $(16 + 16) \times 6400$	58.2% 58.6%	60.7% 61.1%
RMaxTS	non-CoT	4×6400	$55.7\% \pm 0.6\%$	$58.4\% \pm 0.6\%$
		16×6400	57.8%	59.4%
	CoT	4×6400	$56.3\% \pm 0.3\%$	$59.6\% \pm 0.6\%$
		16×6400	59.0%	62.7%
	non-CoT & CoT	$(2 + 2) \times 6400$	$56.1\% \pm 0.8\%$	$60.0\% \pm 0.8\%$
		$(8 + 8) \times 6400$ $(16 + 16) \times 6400$	59.0% 60.2%	63.1% 63.5%

Table 4: A large-scale ablation study to investigate the effectiveness of several algorithmic designs on model training. The results are evaluated on the miniF2F-test dataset.

C ILLUSTRATIVE EXAMPLES OF NON-COT AND COT PROMPTING FOR PROOF COMPLETION

In this section, we present examples of non-CoT and CoT prompting and the generated proofs for `aime_1983_p9`, a problem from the miniF2F benchmark (Zheng et al., 2022). For clarity, inline LaTeX expressions are utilized. Firstly, here is an example of whole-proof generation in non-CoT prompting mode:

Input:

Complete the following Lean 4 code:

```

1056 ``lean4
1057 import Mathlib
1058 import Aesop
1059
1060 set_option maxHeartbeats 0
1061
1062 open BigOperators Real Nat Topology Rat
1063
1064 /-- Show that  $\frac{9x^2 \sin^2 x + 4}{x \sin x} \geq 12$  for  $0 < x < \pi$ . -/
1065 theorem aime_1983_p9 (x : ℝ) (h₀ : 0 < x ∧ x < Real.pi) :
1066   12 ≤ (9 * (x ^ 2 * Real.sin x ^ 2) + 4) / (x * Real.sin x) := by

```

Output:

```

1068   rw [le_div_iff (mul_pos h₀.1 (Real.sin_pos_of_pos_of_lt_pi h₀.1 h₀.2))]
1069   nlinarith [mul_self_nonneg (3 * x * Real.sin x - 2)]
1070   ...

```

Next, here is an example of whole-proof generation in CoT prompting mode. This proof was found using a Monte-Carlo tree search with a truncate-and-resume mechanism (see Section 4 for details). The model receives a prompt including a proof code in progress, ending with a comment block containing the intermediate tactic state from the Lean 4 prover. The subsequent code is generated with guidance from this annotated code prefix.

Input:

Complete the following Lean 4 code with explanatory comments preceding each line of code:

```

1080
1081 ```lean4
1082 import Mathlib
1083 import Aesop
1084
1085 set_option maxHeartbeats 0
1086
1087 open BigOperators Real Nat Topology Rat
1088
1089 /-- Show that  $\frac{9x^2 \sin^2 x + 4}{x \sin x} \geq 12$  for  $0 < x < \pi$ . -/
1090 theorem aime_1983_p9 (x : ℝ) (h₀ : 0 < x ∧ x < Real.pi) :
1091   12 ≤ (9 * (x ^ 2 * Real.sin x ^ 2) + 4) / (x * Real.sin x) := by
1092   /-
1093   To find the minimum value of  $\frac{9x^2 \sin^2 x + 4}{x \sin x}$  for  $0 < x < \pi$ , we need to show
1094   that it is at least 12. We start by noting that the expression can
1095   be rewritten using the division property of inequalities. We then
1096   use the fact that  $\sin x$  and  $x$  are positive in the given range to
1097   establish the necessary inequalities. Finally, we apply these
1098   results to conclude that the minimum value is indeed 12.
1099   -/
1100   -- We start by ensuring that the product  $x * \sin x$  is positive in the
1101   -- given range.
1102   have h₁ : 0 < x * Real.sin x := by
1103     apply mul_pos
1104     --  $x$  is positive in the range  $(0, \pi)$ .
1105     exact h₀.1
1106     --  $\sin x$  is positive in the range  $(0, \pi)$ .
1107     exact Real.sin_pos_of_pos_of_lt_pi h₀.1 h₀.2
1108   -- Using the division property of inequalities, we rewrite the
1109   -- expression.
1110   rw [le_div_iff h₁]
1111   /- tactic state:
1112     x : ℝ
1113     h₀ : 0 < x ∧ x < π
1114     h₁ : 0 < x * x.sin
1115     ⊢ 12 * (x * x.sin) ≤ 9 * (x ^ 2 * x.sin ^ 2) + 4
1116   -/

```

Output:

```

1114   -- This is equivalent to showing that  $9x^2 \sin^2 x - 12x \sin x + 4 \geq 0$ , and the
1115   -- left-hand side can be rewritten as a perfect square  $(3x \sin x - 2)^2$ .
1116   -- We use the fact that  $(3x \sin x - 2)^2$  is non-negative to establish this.
1117   nlinarith [sq_nonneg (3 * x * Real.sin x - 2)]
1118   ```

```

In this case, the proof completion model first extracts the goal expression from the tactic state, *i.e.*, the inequality $12(x \sin x) \leq 9(x^2 \sin^2 x) + 4$ at the last line of the comment block, and then arranges it into the form $9x^2 \sin^2 x - 12x \sin x + 4 \geq 0$ as part of the chain-of-thought process. Finally, the left-hand side is rewritten as a square term to establish non-negativity, closing the proof.

D EXAMPLE SOLUTIONS TO MINIF2F-TEST PROBLEMS

To illustrate the different advantages and complementary of non-CoT mode and CoT mode of prompting, we provide examples where non-CoT mode succeeds while CoT mode fails, and vice versa.

D.1 PROBLEMS THAT NON-COT OUTPERFORMS COT

```

1131 theorem induction_pordlp1on2powklt5on2 (n : ℕ) (h₀ : 0 < n) :
1132   (∏ k in Finset.Icc 1 n, 1 + (1 : ℝ) / 2 ^ k) < 5 / 2 := by

```

```

1134   rw [Finset.Icc]
1135   simp_all [Nat.succ_le_iff, Nat.one_le_iff_ne_zero]
1136   have h1 : 0 < 2 ^ k := by apply pow_pos <|> norm_num
1137   norm_num
1138   have h2 : (2 : ℝ) ≠ 0 := by norm_num
1139   field_simp
1140   rw [div_lt_div_iff]
1141   ring_nf
1142   norm_cast
1143   nlinarith
1144   all_goals norm_cast
1145   all_goals linarith

```

```

1146   theorem imo_1960_p2 (x : ℝ) (h0 : 0 ≤ 1 + 2 * x) (h1 : (1 - Real.sqrt (1
1147     + 2 * x)) ^ 2 ≠ 0)
1148     (h2 : 4 * x ^ 2 / (1 - Real.sqrt (1 + 2 * x)) ^ 2 < 2 * x + 9) : -(1
1149     / 2) ≤ x ∧ x < 45 / 8 := by
1150   norm_num at h0 h1 h2
1151   have h3 : 0 ≤ 1 + 2 * x := by linarith
1152   have h4 : 0 < 1 + Real.sqrt (1 + 2 * x) := by
1153     nlinarith [Real.sqrt_nonneg (1 + 2 * x)]
1154   have h5 : 4 * x ^ 2 / (1 - Real.sqrt (1 + 2 * x)) ^ 2 < 2 * x + 9 :=
1155     by linarith
1156   have h6 : 1 - Real.sqrt (1 + 2 * x) ≠ 0 := by
1157     intro h
1158     apply h1
1159     nlinarith
1160   have h7 : 4 * x ^ 2 / (1 - Real.sqrt (1 + 2 * x)) ^ 2 = (1 + Real.sqrt
1161     (1 + 2 * x)) ^ 2 := by
1162     field_simp [h6]
1163     nlinarith [sq_sqrt (show 0 ≤ 1 + 2 * x by linarith)]
1164   rw [h7] at h5
1165   constructor <|> nlinarith [sq_sqrt (show 0 ≤ 1 + 2 * x by linarith)]

```

D.2 PROBLEMS THAT CoT OUTPERFORMS NON-CoT

```

1166   theorem mathd_algebra_459 (a b c d : ℚ) (h0 : 3 * a = b + c + d) (h1 : 4
1167     * b = a + c + d)
1168     (h2 : 2 * c = a + b + d) (h3 : 8 * a + 10 * b + 6 * c = 24) : ↑d.den +
1169     d.num = 28 := by
1170   /-
1171   Given the system of equations involving the times that Anna (a),
1172   Bertram (b), Carli (c), and David (d) can hold their breaths:
1173   1. 3a = b + c + d
1174   2. 4b = a + c + d
1175   3. 2c = a + b + d
1176   4. 8a + 10b + 6c = 24
1177   We need to find the sum of the numerator and the denominator of the
1178   fraction representing David's time in minutes, simplified to its
1179   lowest terms, and show that it is 28.
1180   First, we solve the system of equations step by step:
1181   - From the first equation, we can express d in terms of a, b, and c:
1182     d = 3a - b - c
1183   - Substitute d into the second equation:
1184     4b = a + c + (3a - b - c)
1185     Simplifying, we get:
1186     4b = 4a - b
1187     5b = 4a
1188     b = 4a/5
1189   - Substitute b into the third equation:
1190     2c = a + 4a/5 + (3a - b - c)
1191     Simplifying, we get:
1192     2c = a + 4a/5 + 3a - 4a/5 - c

```

```

1188      2c = 4a - c
1189      3c = 4a
1190      c =  $\frac{4a}{3}$ 
1191      - Substitute b and c into the fourth equation:
1192       $8a + 10\left(\frac{4a}{5}\right) + 6\left(\frac{4a}{3}\right) = 24$ 
1193      Simplifying, we get:
1194       $8a + 8a + 8a = 24$ 
1195       $24a = 24$ 
1196       $a = 1$ 
1197      Now, we find b and c:
1198      -  $b = \frac{4a}{5} = \frac{4}{5}$ 
1199      -  $c = \frac{4a}{3} = \frac{4}{3}$ 
1200      Finally, we find d:
1201      -  $d = 3a - b - c = 3 \cdot 1 - \frac{4}{5} - \frac{4}{3} = 3 - \frac{4}{5} - \frac{4}{3} = \frac{45}{15} - \frac{12}{15} - \frac{20}{15} = \frac{13}{15}$ 
1202      The fraction  $\frac{13}{15}$  is already in its simplest form, and the sum of the
1203      numerator and the denominator is:
1204       $13 + 15 = 28$ 
1205      -/
1206      have h4 : a = 1 := by
1207      have h5 : 8 * a + 10 * b + 6 * c = 24 := h3
1208      have h6 : 24 * a = 24 := by
1209      -- Solve the system of equations to find the value of a
1210      linarith
1211      exact mul_right_cancel0 (show (24 : ℚ) ≠ 0 by norm_num)
1212      (by linarith)
1213      have h5 : b = 4 / 5 := by
1214      -- Solve for b using the value of a
1215      linarith
1216      have h6 : c = 4 / 3 := by
1217      -- Solve for c using the value of a
1218      linarith
1219      have h7 : d = 13 / 15 := by
1220      -- Solve for d using the values of a, b, and c
1221      linarith
1222      -- Sum the numerator and the denominator of d
1223      norm_num [h7]
1224      <.> rfl
1225
1226      theorem numbertheory_x5neqy2p4 (x y : ℤ) : x ^ 5 ≠ y ^ 2 + 4 := by
1227      /-
1228      To show that for any two integers x and y,  $x^5 \neq y^2 + 4$ , we can use
1229      modular arithmetic to demonstrate that the equation  $x^5 \equiv y^2 + 4 \pmod{11}$ 
1230      has no solutions. Specifically, we will consider the possible values
1231      of x and y modulo 11 and show that none of these combinations satisfy
1232      the equation.
1233      -/
1234      -- Use modular arithmetic to show that the equation has no solutions
1235      modulo 11.
1236      have : ¬(x ^ 5 % 11 = (y ^ 2 + 4) % 11) := by
1237      -- Consider all possible values of x modulo 11.
1238      have : x % 11 = 0 ∨ x % 11 = 1 ∨ x % 11 = 2 ∨ x % 11 = 3 ∨ x % 11 =
1239      4 ∨ x % 11 = 5 ∨
1240      x % 11 = 6 ∨ x % 11 = 7 ∨ x % 11 = 8 ∨ x % 11 = 9 ∨ x % 11 = 10
1241      := by
1242      omega
1243      -- Consider all possible values of y modulo 11.
1244      have : y % 11 = 0 ∨ y % 11 = 1 ∨ y % 11 = 2 ∨ y % 11 = 3 ∨ y % 11 =
1245      4 ∨ y % 11 = 5 ∨
1246      y % 11 = 6 ∨ y % 11 = 7 ∨ y % 11 = 8 ∨ y % 11 = 9 ∨ y % 11 = 10
1247      := by
1248      omega
1249      -- For each combination of x and y modulo 11, check if the equation
1250      holds.
1251      rcases this with (h | h | h | h | h | h | h | h | h | h) <.>

```

```

1242   rcases this with (h' | h' | h' | h' | h' | h' | h' | h' | h' | h' | h')
1243   <.>
1244   simp [h, h', pow_succ, Int.mul_emod, Int.add_emod]
1245   -- Use the result from modular arithmetic to conclude the theorem.
1246   omega
1247
1248 theorem amc12_2000_p12 (a m c : ℕ) (h₀ : a + m + c = 12) :
1249   a * m * c + a * m + m * c + a * c ≤ 112 := by
1250   /-
1251   Given nonnegative integers  $A, M,$  and  $C$  such that  $A + M + C = 12$ , we need
1252   to find the maximum value of  $A \cdot M \cdot C + A \cdot M + M \cdot C + A \cdot C$ .
1253   To solve this, we can use the fact that the expression
1254    $A \cdot M \cdot C + A \cdot M + M \cdot C + A \cdot C$  is a polynomial in terms of  $A, M,$  and  $C$ .
1255   We can try different values of  $A, M,$  and  $C$  that sum to 12 and compute
1256   the value of the polynomial for each combination. The maximum value
1257   of the polynomial will then be the solution.
1258   -/
1259   -- We use a lemma that states if the sum of three numbers is less than
1260   -- or equal to a certain value, then the product of these numbers is
1261   -- also less than or equal to a certain value.
1262   have h₁ : a ≤ 12 ∧ m ≤ 12 ∧ c ≤ 12 := by
1263     constructor
1264     · -- Since  $a + m + c = 12$ ,  $a$  cannot be greater than 12.
1265       omega
1266     · constructor
1267       · -- Similarly,  $m$  cannot be greater than 12.
1268         omega
1269       · -- And  $c$  cannot be greater than 12.
1270         omega
1271   -- Using the lemma, we compute the value of the polynomial for
1272   -- different values of  $a, m,$  and  $c$  that sum to 12.
1273   cases' h₁ with h₁_left h₁_right
1274   cases' h₁_right with h₁_middle h₁_right
1275   -- We use interval_cases to try different values of  $a, m,$  and  $c$ .
1276   interval_cases a <.> interval_cases m <.> interval_cases c <.>
1277   -- For each combination, we compute the value of the polynomial and
1278   -- check if it is less than or equal to 112.
1279   simp_all (config := {decide := true})
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295

```