

UCO: A Multi-Turn Interactive Reinforcement Learning Method for Adaptive Teaching with Large Language Models

Shouang Wei¹, Min Zhang^{1,†,‡}, Xin Lin^{1,†}, Bo Jiang¹, Kun Kuang²,
Jingyuan Chen², Zhongxiang Dai³

¹East China Normal University,

²Zhejiang University, ³The Chinese University of Hong Kong, Shenzhen

[‡]Project leader 52285901019@stu.ecnu.edu.cn [†]Corresponding authors: mzhang@cs.ecnu.edu.cn, xlin@cs.ecnu.edu.cn

Abstract

Large language models (LLMs) are shifting from answer providers to intelligent tutors in educational settings. However, current supervised fine-tuning methods only learn surface-level teaching patterns without dynamic adaptation. Recent reinforcement learning approaches address this, yet two key challenges remain. First, they evaluate teaching effectiveness solely based on whether students produce correct outputs, unable to distinguish whether students genuinely understand or echo teacher-provided answers. Second, they cannot perceive students' evolving cognitive states through dialogue, failing to adapt teaching strategies accordingly. We propose Unidirectional Cognitive Optimization (UCO) to address these challenges. UCO employs a multi-turn interactive reinforcement learning paradigm with two synergistic reward functions: the Progress Reward captures cognitive advancement, evaluating students' transition from confusion to comprehension, and the Scaffold Reward dynamically identifies each student's Zone of Proximal Development (ZPD), encouraging the teacher to conduct productive teaching within this zone. We evaluate UCO against 11 baseline models on BigMath and MathTutorBench. Results demonstrate that UCO outperforms all equivalent-scale models and achieves performance comparable to leading closed-source models. Code is available at <https://github.com/ECNU-RAIL/UCO-EMNLP2026>.

1 Introduction

Large language models are shifting from "answer providers" to "intelligent tutors" in educational applications. Mainstream methods rely on supervised fine-tuning (SFT) with human-annotated or model-generated teaching dialogues (Macina et al., 2023a; Wang et al., 2024c; Liu et al., 2024c), showing how teachers guide students through questioning, prompting, and feedback. However, this imitation




Figure 1: Two Teaching Methods: Comparison Between Non-Adaptive (a) and Adaptive (b) approaches.

approach faces a key challenge: SFT-trained models repeat fixed patterns from training data, lacking the dynamic adaptability of real teachers.

Current research explores reinforcement learning (RL) methods to break through this bottleneck (Fahid et al., 2024; Riedmann et al., 2025; Dinucu-Jianu et al., 2025; Qu et al., 2026). Early work primarily adopts single-turn RL paradigms (Zhou et al., 2024; Jia et al., 2026a; Poesia et al., 2024; Sheng et al., 2025; Hong et al., 2026; Kong et al., 2026). However, real teaching is inherently a multi-turn dynamic process. Thus, recent research shifts to multi-turn RL methods (Dinucu-Jianu et al., 2025; Xiong et al.; Li et al., 2025), enabling iterative teacher-student interactions. Nevertheless, multi-turn RL approaches primarily rely on the correctness of student outputs as evaluation criteria, leading to two critical deficiencies. First, relying solely on answer correctness as the reward signal, when students output correct

answers, models cannot distinguish superficial correctness from genuine understanding (students may merely guess correctly or repeat answers that the teacher revealed), and consequently cannot identify which teaching steps truly promote cognitive improvement. Second, since the current reward function design prevents models from perceiving students' actual understanding levels, models cannot select teaching strategies that match students' comprehension. As illustrated in Figure 1(a), this non-adaptive teaching fails to identify student cognitive states and provide cognitive-matched strategies, leading to teaching failure.

To effectively address these limitations, we propose the **Unidirectional Cognitive Optimization (UCO)** method. UCO adopts a multi-turn interactive reinforcement learning paradigm. The teacher and student models interact continuously to generate online training rollouts dynamically. The approach computes rollout returns based on reward mechanisms and uses GRPO (Shao et al., 2024) to update the teacher policy. As shown in Figure 1(b), our adaptive teaching approach identifies student cognitive states and provides matched teaching strategies. This method ultimately achieves teaching success. The core innovation lies in two specifically designed reward functions:

- **Progress Reward.** From an information theory perspective, we formalize student cognitive progress as an entropy reduction process (Press et al., 2024; Davis et al., 2007). Students transition from high-entropy states with high uncertainty to low-entropy states with high certainty. However, computing the entropy of student cognitive states is impractical. We argue that increases in student model probability for outputting correct answers essentially equal cognitive state transitions. Based on this insight, we design dual-dimensional indicators to quantify this entropy reduction process. We combine these two dimensions with weighted aggregation to provide reward signals for each interaction turn.

1. **Potential Capability Score:** For each teacher output, we first use an oracle model to generate multiple candidate correct student responses. We then compute the log-probability of the student model for generating each candidate response. Higher values indicate the student model

has high confidence in the correct response. This corresponds to cognitive entropy reduction. Lower values reflect student confusion states. We apply the hyperbolic tangent function to map log-probabilities to the $[-1, 1]$, yielding the Potential Capability Score.

2. **Semantic Quality Score:** We use text embedding models to compute maximum cosine similarity between student outputs and candidate correct answers. We designate this similarity measure as the Semantic Quality Score.

- **Scaffold Reward.** Pedagogical practice reveals that overly simplistic tasks induce boredom, whereas prohibitively challenging ones elicit frustration. Difficulty should remain within the student's "productive struggle" range, termed by Vygotsky as the Zone of Proximal Development (ZPD) (Shabani et al., 2010). To translate this into a computable mechanism, we divide teacher behaviors by cognitive load into five ordered levels, from high to low: Metacognitive, Strategic, Conceptual, Step-by-Step, and Example Demonstration. Leveraging this tiered scaffolding framework, we operationalize a mechanism for adaptive ZPD demarcation. We first evaluate the student model's success probability under each scaffold level and identify the highest success level as mastery ability. We then reduce scaffolding by one level, moving tasks into the ZPD region requiring moderate productive struggle. The system provides positive rewards when the selected scaffold falls within the ZPD, and applies proportionally increasing penalties for deviations.

By combining these reward functions via weighted aggregation, the method computes rewards after each turn. Then we update the teacher model's policy for cognition-oriented adaptive teaching optimization. Our key contributions are:

- We propose a novel unidirectional cognitive optimization method that generates training trajectories online through iterative teacher-student interactions, dynamically adjusting teaching strategies as student cognitive states evolve, enabling truly adaptive instruction.
- We design a cognition-oriented reward function with two components: a progress re-




Figure 2: Overview of the Unidirectional Cognitive Optimization (UCO) method, which optimizes the teacher model through multi-turn interactions with the student model, designing progress reward and scaffold reward to evaluate teacher outputs at each turn, and updating the teacher model via GRPO.

ward quantifying cognitive improvement, and a ZPD-based scaffolding reward. These rewards synergistically ensure guidance aligned with students’ cognitive levels.

- We compare our method against 11 baselines on BigMath and MathTutorBench. Our UCO model surpasses all comparable-sized models and achieves performance comparable to advanced closed-source models.

2 Related Work

Math Tutoring with Large Language Models. Effective teaching builds student cognition through scaffolding guidance rather than direct answers (Van de Pol et al., 2010; Bliss et al., 1996; Rosenshine et al., 1992). However, LLMs center their training on answer generation, fundamentally conflicting with pedagogical principles (Tack and Piech, 2022; Kucheria et al., 2025; Puech et al., 2024; Macina et al., 2023b; Zhang et al., 2025a). Transforming LLMs into qualified teaching partners becomes an active research direction. Mathematics serves as an ideal testbed due to its logical rigor and quantifiable nature (Luo et al.; Liu et al., 2024b; Jia et al., 2026b). Early research constructs structured educational dialogue data through multi-agent simulation (Liu et al., 2024c; De Marco et al., 2024), establishing the foundation for learning teaching skills. Researchers then build RL-specific benchmarks like Big-Math (Albalak et al., 2025) and MathTutorBench (Macina et al., 2025), advancing

research standardization. To enable dynamic optimization of teaching strategies, researchers adopt RL as the mainstream framework. One approach introduces offline preference learning to improve teaching content accuracy (Xiong et al.; Chu et al., 2025), optimizing problem-solving trajectory quality. Yet this offline nature prevents real-time strategy adjustment based on student feedback, lacking dynamic interaction capabilities. Subsequent research builds online interactive environments with simulated students (Dinucu-Jianu et al., 2025), directly rewarding specific teaching behaviors. However, they rely on simple behavioral rewards, failing to measure the long-term goal of student understanding. We design an online RL framework to overcome these limitations. Our teacher model optimizes by real-time perceiving student cognitive states through interaction, achieving truly adaptive teaching in authentic interactive contexts.

Reward Mechanisms for Teaching Dialogues.

Reinforcement learning formulates teaching dialogues as sequential decision-making processes (Klissarov et al.; Wang et al., 2024a), optimizing pedagogical strategies through reward signals. The reward function architecture serves as the primary determinant of system performance. Early work employs outcome-oriented sparse rewards (Peng et al., 2025; Liao et al., 2024; Zhong et al., 2024), providing feedback solely at dialogue termination based on response accuracy. This faces severe credit assignment problems as multi-turn di-

alogues make identifying critical teaching actions difficult, leading to inefficient policy learning. Subsequent research shifts to process-oriented dense rewards (Zhang et al., 2025c; Yin et al., 2025; Li and Li), evaluating teaching effectiveness after each dialogue turn. However, these mechanisms suffer from goal alienation by focusing only on student output correctness while ignoring whether students understand the content. This encourages teacher models to take shortcuts, receiving high rewards whenever student models output correct steps, regardless of genuine understanding. We propose a new reward function to address these limitations by analyzing student model cognitive state changes to evaluate cognitive progress and measuring how teaching interventions match the student’s proximal development zone. This achieves a paradigm shift from answer correctness to understanding depth.

3 Method

3.1 Problem Formalization

We formalize adaptive teaching dialogue generation as a POMDP $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{R}, \gamma)$, where $\mathcal{S}$ contains problem descriptions and dialogue history, $\mathcal{A}$ represents pedagogical strategies, $\mathcal{O}$ corresponds to student responses, $\mathcal{T}$ models state transitions, $\mathcal{R}$ evaluates teaching effectiveness, and γ balances immediate and future rewards.

As illustrated in Figure 2, we propose Unidirectional Cognitive Optimization (UCO) to train an optimal teacher model π_{teacher}^* that adaptively guides students to solve mathematical problems. Specifically, UCO tightly couples three core components into a joint optimization loop: (1) multi-turn interactive rollout generation driven by the teacher and student models, (2) cognitive state evaluation via a dual-objective reward function assisted by an oracle model, and (3) teacher policy updates using GRPO. We aim to maximize the expected cumulative reward across dialogue interactions:

$$\pi_{\text{teacher}}^* = \arg \max_{\pi_{\text{teacher}}} \mathbb{E}_{q \sim \mathcal{D}, \tau \sim \rho} \left[\sum_{t=0}^T \gamma^t r_t^{(i)} \right], \quad (1)$$

where $\pi_{\text{student}}^{\text{fixed}}$ is the fixed student model, q is a problem sampled from distribution $\mathcal{D}$, τ is a dialogue rollout sampled from $\rho = \rho(\pi_{\text{teacher}}, \pi_{\text{student}}^{\text{fixed}})$, which denotes the distribution of interactions between the teacher and student models, and T is the maximum dialogue turns. The state at turn t for the

i -th rollout is $s_t^{(i)} = \{q, a_0^{(i)}, o_1^{(i)}, \dots, a_{t-1}^{(i)}, o_t^{(i)}\}$, comprising the problem and dialogue history up to turn t , with reward $r_t^{(i)} = \mathcal{R}(s_t^{(i)}, a_t^{(i)})$.

3.2 Dual-Objective Reward Function

We design a dual-objective reward function to evaluate teaching effectiveness, comprising the progress reward and the scaffold reward.

3.2.1 Progress Reward

Based on Shannon’s information theory, effective teaching should maximize information gain, guiding students from uncertainty to certainty. This shift from high entropy to low entropy is termed **Cognitive Progress**. Formally, given a student’s cognitive state ψ_t at time t , cognitive progress after action a_t is measured by entropy reduction:

$$\Delta H_t = H(\psi_t) - H(\psi_{t+1}|a_t), \quad (2)$$

where $H(\cdot)$ denotes information entropy.

However, directly calculating $H(\psi_t)$ is impractical. We observe that increases in correct output probability essentially reflect transitions from high-entropy to low-entropy states. Based on this, we design r_{progress} to approximate cognitive progress through two dimensions: the student’s **potential capacity** and **semantic quality** of outputs.

Potential Capability Score. This term directly measures the confidence of the fixed student model π_S^{fixed} in correct paths. We use it as a core proxy for cognitive progress. We calculate it as:

$$f_{\text{potential}} = \tanh \left(\alpha \cdot \max_{k=1}^n \sum_{j=1}^{m_k} \log \pi_S^{\text{fixed}}(c_{t,j}^k | s_t, a_t, c_{t,1:j-1}^k) \right), \quad (3)$$

where $\{c_t^1, c_t^2, \dots, c_t^n\}$ is the set of candidate correct responses generated by the oracle model for the student model in the t -th interaction round. Each candidate c_t^k is a token sequence of length m_k , where $c_{t,j}^k$ denotes the j -th token in candidate k . The term $\max_{k=1}^n$ captures the highest log probability that the student model assigns across all oracle-generated candidates, computed autoregressively over the token sequence. This reflects the student model’s peak performance potential at the current step. A higher probability value indicates that the student model has greater confidence in generating correct reasoning paths. This corresponds to a

cognitive shift from uncertainty (high entropy) to certainty (low entropy), serving as a key signal for measuring cognitive progress.

We introduce the tanh function and temperature coefficient α to convert log-likelihoods. The range of log-likelihoods is $(-\infty, 0]$. The tanh function maps values to the bounded range $[-1, 1]$. This ensures numerical stability and penalizes outputs with insufficient confidence. These outputs still reflect the student’s confused state. The temperature coefficient $\alpha > 0$ controls reward sensitivity. Through this design, $f_{\text{potential}}$ quantifies the student model’s cognitive progress.

Semantic Quality Score. Besides potential ability assessment, we also need to measure whether the student’s actual output is semantically close to the correct answer. We define this component as:

$$f_{\text{semantic}} = \max_{c_j \in \mathcal{C}} \text{sim}(\mathbf{e}(o_{t+1}), \mathbf{e}(c_j)) - \delta, \quad (4)$$

where $\mathbf{e}(\cdot)$ denotes a text embedding model (e.g., BAAI General Embedding (Chen et al., 2024)) that maps text into a vector space, $\text{sim}(\cdot, \cdot)$ measures cosine similarity, and $\mathcal{C}$ is the set of candidate correct answers. The max operator captures that semantic alignment with any correct answer indicates concept mastery. The bias term δ filters out semantically ambiguous or low-quality outputs.

We integrate potential ability and semantic quality through the hyperparameter $\lambda \in [0, 1]$, capturing both internal ability improvement and external semantic performance to provide a comprehensive cognitive progress signal:

$$r_{\text{progress}} = \lambda f_{\text{potential}} + (1 - \lambda) f_{\text{semantic}}. \quad (5)$$

3.2.2 Scaffold Reward

The scaffold reward r_{scaffold} ensures that teaching difficulty is calibrated to students’ current abilities. Tasks that are too easy cause students to lose interest, while tasks that are too hard lead to frustration. Ideal teaching maintains task difficulty within a range where students can succeed through moderate effort, a principle Vygotsky termed the Zone of Proximal Development. We transform this theory into a computable mechanism by discretizing teaching actions into five **scaffold levels** $\mathcal{L} = \{\ell_0, \dots, \ell_4\}$ based on cognitive load from high to low: ℓ_0 (**Metacognitive Hints**) uses open-ended questions to guide student thinking without providing substantive hints; ℓ_1 (**Strategic Hints**) provides macro-level directions without specific meth-

ods; ℓ_2 (**Conceptual Hints**) identifies key knowledge points or theorems; ℓ_3 (**Step-by-Step Hints**) provides specific operational instructions; and ℓ_4 (**Example Demonstration Hints**) offers complete solutions to similar problems.

We design a strategy for dynamically identifying the ZPD according to this hierarchy. In each round, the oracle model generates five hints $\{\text{hint}_{\ell_0}, \dots, \text{hint}_{\ell_4}\}$ for the student model. We evaluate the probability that the student generates correct responses under each hint:

$$P(\ell) = \log \pi_S(c_t | s_t, \text{hint}_\ell), \quad (6)$$

where c_t represents a sample from the candidate response set, s_t denotes the current student state, and hint_ℓ is the hint at level ℓ .

We identify the level that maximizes $P(\ell)$ to determine the student’s ability range, then lower the target scaffold level by one grade, shifting teaching difficulty from the comfort zone to the ZPD:

$$\ell_{\text{ZPD}} = \max(0, \arg \max_{\ell \in \mathcal{L}} P(\ell) - 1), \quad (7)$$

where the minus-one operation sets the target one level below, creating a moderate challenge. Based on ℓ_{ZPD} , we define the scaffolding reward:

$$r_{\text{scaffold}} = \begin{cases} \sigma(P(\ell(a_t))) + 0.5, & \ell(a_t) = \ell_{\text{ZPD}} \\ -c|\Delta\ell_t|, & \text{otherwise} \end{cases} \quad (8)$$

where $\text{index}(\cdot)$ converts levels to integers (0 to 4), $P(\ell(a_t))$ is the log probability of correct response, σ is the Sigmoid function, $c > 0$ is a penalty coefficient, and $\Delta\ell_t = \text{index}(\ell(a_t)) - \ell_{\text{ZPD}}$ denotes the level distance. When $\ell(a_t) = \ell_{\text{ZPD}}$, the model receives 0.5 plus a dynamic reward $\sigma(P(\ell(a_t)))$. Otherwise, it receives a penalty proportional to $|\Delta\ell_t|$. This mechanism enables the teacher to deliver precisely calibrated scaffolding within the ZPD, mitigating both over- and under-scaffolding. The total reward combines both components:

$$r_t^{(i)} = r_{\text{progress}} + r_{\text{scaffold}}, \quad (9)$$

where $r_t^{(i)}$ is the reward at round t in rollout i .

3.3 Policy Optimization via GRPO

We use GRPO to jointly optimize the teacher model π_{teacher} through a continuous feedback loop among the three UCO components. During each training iteration, the components interact as follows: First, the teacher model interacts with the fixed

student model $\pi_{\text{student}}^{\text{fixed}}$ to generate G diverse pedagogical trajectories for each problem $q \in \mathcal{B}$. Second, the oracle model is introduced to evaluate the student’s hidden cognitive state turn-by-turn, computing $r_t^{(i)}$ to form dense feedback. Finally, these rewards are aggregated into a cumulative return $\mathcal{R}_i = \sum_{t=0}^T \gamma^t r_t^{(i)}$ for the i -th rollout, where $\gamma \in [0, 1]$ is the discount factor. The intra-group advantage is then computed as $A_i = (\mathcal{R}_i - \mu_q) / (\sigma_q + \epsilon)$, with μ_q and σ_q denoting the mean and standard deviation of returns within the same group. The advantage A_i guides the policy gradient update, with a KL divergence term preventing deviations from the reference policy $\pi_{\text{teacher}}^{\text{ref}}$. We define the final GRPO objective $\mathcal{J}$ as:

$$\mathcal{J} = \mathbb{E}_{q \sim \mathcal{B}} \left[\frac{1}{G} \sum_{i=1}^G \left(A_i \cdot \sum_{t=0}^T \log \pi_{\text{teacher}}(a_t^{(i)} | s_t^{(i)}) - \xi \cdot D_{\text{KL}}[\pi_{\text{teacher}}(\cdot | s_t^{(i)}) \| \pi_{\text{teacher}}^{\text{ref}}(\cdot | s_t^{(i)})] \right) \right], \quad (10)$$

where $\xi > 0$ is the KL penalty weight. The updated teacher policy is deployed in the next rollout phase, forming closed-loop joint optimization.

4 Experiments

4.1 Experimental Setup

Dataset. We train on 2,000 randomly sampled BigMath instances (Albalak et al., 2025), following PedagogicalRL (Dinucu-Jianu et al., 2025), to both validate our method and reduce RL training cost.

Evaluation. We evaluate tutoring performance on two complementary benchmarks. On BigMath, we follow the exact evaluation protocol of PedagogicalRL, using the same 500 held-out test questions, the same student model (Meta-Llama-3.1-8B-Instruct-FP8 (Grattafiori et al., 2024)), and the same judge model (Gemma-3-27B-IT (Kamath et al., 2025)), replacing only the teacher model with ours. We measure: (1) Δ Solve Rate, quantifying student problem-solving improvement; (2) Leak Solution, detecting explicit solution divulgence via LLM adjudication; and (3) Ped-RM micro/macro scores from MathTutorBench (Macina et al., 2025), assessing pedagogical quality. On MathTutorBench, we follow the evaluation protocol in the original paper and assess nine tasks across three domains: *Math Expertise* (Problem Solving accuracy, Socratic Questioning BLEU), *Student*

Understanding (Solution Correctness F1, Mistake Location F1, Mistake Correction accuracy), and *Pedagogy* (win rates on Scaffolding Generation, Pedagogical IF, and Hard variants in dialogues).

Baselines. We compare against 11 models: three closed-source systems (GPT-4o (Hurst et al., 2024), LearnLM 1.5 Pro and 2.0 Flash (Team et al., 2024)), four open-source models spanning 7B-671B parameters (Qwen2.5 variants (Team, 2024), DeepSeek-V3 (Liu et al., 2024a)), and four education-specialized models (Qwen2.5-7B-SFT (Macina et al., 2023a), Qwen2.5-7B-MDPO (Xiong et al.), SocraticLM (Liu et al., 2024c), Qwen2.5-7B-RL (Dinucu-Jianu et al., 2025)).

Training Setting. We use Qwen2.5-7B-Instruct (Team, 2024) as π_{teacher} , Qwen2.5-14B-Instruct (Team, 2024) as $\pi_{\text{student}}^{\text{fixed}}$, and Gemini-2.5-Pro-Exp (Comanici et al., 2025) for candidate generation (1 teacher response per scaffolding level, $K = 5$ student responses per turn). Complete details appear in Appendix B.

4.2 Main Results

Performance on BigMath Benchmark. Table 1 shows that existing models struggle to balance three dimensions: teaching effectiveness (Δ Solve Rate), answer leakage (Leak Solution), and teaching quality (Ped-RM micro/macro).

General models prioritize answers over teaching. Closed-source GPT-4o achieves 33.1% Δ Solve Rate but has 35.2% Leak Solution and only 1.5/−0.3 Ped-RM scores. Among open-source models, larger parameters worsen the imbalance. The Qwen2.5 series increases Δ Solve Rate from 11.3% to 38.7% as parameters grow from 7B to 72B, while Leak Solution rises from 29.3% to 61.0% and Ped-RM stays negative. DeepSeek-V3 achieves the highest Δ Solve Rate of 39.3% but shows 46.6% Leak Solution with −1.5/−0.8 Ped-RM scores, revealing large-scale models fundamentally function as “solvers” rather than teachers.

Education-specialized models exhibit a different pattern. SFT and MDPO achieve lower Leak Solution rates but their Δ Solve Rates drop to 8.9% and 16.4%. SocraticLM and Qwen2.5-7B-RL achieve better balance, with Qwen2.5-7B-RL reaching 29.1% Δ Solve Rate and 15.1% Leak Solution, though its Ped-RM scores of 3.6/3.1 leave room for improvement. LearnLM 2.0 reaches the best Ped-RM scores of 6.8/6.4 and lowest Leak Solution of 0.9%, but its Δ Solve Rate is only 4.3%,

demonstrating the inherent trade-off: strict adherence to pedagogical principles can compromise student learning outcomes.

Models	Δ Solve Rate (%) ↑	Leak Solution (%) ↓	Ped-RM micro/macro ↑
<i>Closed-source</i>			
GPT-4o	33.1	35.2	1.5/-0.3
LearnLM 1.5 Pro	1.5	2.6	5.9/5.3
LearnLM 2.0 Flash	4.3	0.9	6.8/6.4
<i>Open-source</i>			
Qwen2.5-7B	11.3	29.3	-0.2/-0.5
Qwen2.5-14B	29.3	41.9	-0.6/-1.2
Qwen2.5-72B	38.7	61.0	1.8/-0.4
DeepSeek-V3	39.3	46.6	-1.5/-0.8
<i>Education-specific</i>			
Qwen2.5-7B-SFT	8.9	36.0	-0.3/-0.7
Qwen2.5-7B-MDPO	16.4	35.6	0.2/-0.3
SocraticLM	15.9	40.4	1.7/1.7
Qwen2.5-7B-RL	29.1	15.1	3.6/3.1
UCO	30.2	12.9	4.6/4.5

Table 1: Performance on BigMath benchmark. Results for all models except UCO are from the PedagogicalRL.

UCO achieves synergistic optimization with only 7B parameters. It reaches 30.2% Δ Solve Rate, surpassing similar-sized models and approaching larger models. UCO achieves 12.9% Leak Solution, representing a 14.6 percentage point improvement over the strongest teaching baseline Qwen2.5-7B-RL, and 4.6/4.5 Ped-RM scores (28% improvement). This three-way balance validates our dual-objective reward mechanism.

Our core advantage lies in modeling student cognitive states. Learning progress reward quantifies cognitive changes, providing rewards only when students truly understand reasoning steps, thus improving Δ Solve Rate. Scaffolding reward operationalizes ZPD theory by estimating student ability across five levels, identifying the zone of proximal development, and encouraging adaptive strategies that avoid direct answers while preventing inefficient over-questioning. UCO leverages these differentiable cognitive signals, surpassing similarly sized baselines across all three dimensions and achieving comprehensive performance that rivals closed-source models.

Comprehensive Evaluation on MathTutorBench. Table 2 shows evaluation results on MathTutorBench. We compare UCO with education-specific models across nine tasks in three categories: Math Expertise, Student Understanding, and Pedagogy. Qwen2.5-7B-SFT performs well in Math Expertise but lags significantly in Student Understanding and Pedagogy. Qwen2.5-7B-MDPO improves math

solving and student understanding through direct preference optimization, but causes severe Pedagogy degradation. SocraticLM excels in Pedagogy but shows weaknesses in Math Expertise and Student Understanding. Qwen2.5-7B-RL achieves balance across dimensions but lacks overall strength. In contrast, UCO performs strongly in Math Expertise and Student Understanding while achieving optimal Pedagogy performance. This stems from UCO’s cognitive state modeling mechanism, where learning progress rewards track student cognitive changes and scaffolding rewards dynamically adjust teaching strategies, maintaining high-quality teaching in complex multi-turn interactions.

4.3 Ablation Study

Ablation Study on BigMath Dataset. Table 4 validates the necessity of our reward components through ablation studies on BigMath. The entire UCO model achieves 30.2% Δ Solve Rate, 12.9% Leak Solution, and 4.6/4.5 Ped-RM micro/macro. Removing scaffolding rewards causes the most severe degradation. Compared to the entire setup, Δ Solve Rate decreases by 6.7 percentage points to 23.5%, Leak Solution rises by 7.3 percentage points to 20.2%, and Ped-RM micro/macro drops to 2.2/2.3. Without scaffolding guidance, the model fails to adapt to students’ zones of proximal development, resulting in inefficient teaching. Removing progress rewards leads to smaller but notable declines. Compared to the entire setup, Δ Solve Rate drops by 1.6 percentage points to 28.6%, Leak Solution increases by 2.7 percentage points to 15.6%, and Ped-RM micro/macro decreases to 4.1/3.9. The model struggles to track student understanding and cannot adjust teaching strategies based on learning states. All reward configurations outperform the backbone, confirming the effectiveness of our reward design.

Ablation Study on MathTutorBench Dataset. Table 3 further validates our reward function through ablation studies on MathTutorBench. The entire UCO model performs well across all three dimensions. We remove scaffolding rewards and observe significant declines: 3.4% in Math Expertise, 9.0% in Student Understanding, and 10.6% in Pedagogy. Without scaffolding guidance, the model fails to break down content into progressive steps. It causes a mismatch between teaching strategies and student capabilities. Removing progress rewards has a greater impact on recognizing stu-

Model	Math Expertise		Student Understanding			Pedagogy			
	Problem solving	Socratic questioning	Solution correctness	Mistake location	Mistake correction	Teacher response generation			
	acc.	BLEU	F1	micro F1	acc.	scaff.	ped.IF	scaff. [hard]	ped.IF [hard]
<i>Education-specific Models</i>									
Qwen2.5-7B-SFT	0.77	0.24	0.27	0.45	0.10	0.64	0.58	0.57	0.59
Qwen2.5-7B-MDPO	0.86	0.23	0.62	0.39	0.03	0.37	0.60	0.47	0.56
SocraticLM	0.73	0.32	0.05	0.39	0.23	0.39	0.39	0.28	0.28
Qwen2.5-7B-RL	0.83	0.23	0.67	0.35	0.05	0.57	0.72	0.61	0.69
UCO	0.89	0.24	0.65	0.45	0.08	0.68	0.74	0.72	0.77

Table 2: Evaluation on MathTutorBench across math expertise, student understanding, and pedagogy.

Model	Math Expertise		Student Understanding			Pedagogy			
	Problem solving	Socratic questioning	Solution correctness	Mistake location	Mistake correction	Teacher response generation			
	acc.	BLEU	F1	micro F1	acc.	scaff.	ped.IF	scaff. [hard]	ped.IF [hard]
<i>Backbone Model</i>									
Qwen2.5-7B-Instruct	0.87	0.23	0.63	0.39	0.04	0.37	0.60	0.45	0.56
UCO (full)	0.89	0.24	0.65	0.45	0.08	0.68	0.74	0.72	0.77
UCO w/o scaffold reward	0.86	0.23	0.61	0.40	0.05	0.56	0.70	0.67	0.71
UCO w/o progress reward	0.88	0.24	0.63	0.42	0.07	0.61	0.71	0.69	0.74

Table 3: Ablation study on MathTutorBench showing impact of reward components across evaluation dimensions.

Models	Δ Solve Rate (%) ↑	Leak Solution (%) ↓	Ped-RM micro/macro ↑
<i>Backbone Model</i>			
Qwen2.5-7B-Instruct	11.3	29.3	-0.2/-0.5
UCO (full)	30.2	12.9	4.6/4.5
UCO w/o scaffold reward	23.5	20.2	2.2/2.3
UCO w/o progress reward	28.6	15.6	4.1/3.9

Table 4: Ablation study on BigMath on the contribution of progress and scaffolding reward components.

dent understanding states (Student Understanding) and dynamically adjusting teaching strategies (Pedagogy), but less impact on accurately delivering mathematical knowledge (Math Expertise). It indicates that progress rewards primarily regulate teaching pace based on student feedback.

Additional Ablation Studies. To thoroughly validate UCO, Appendix C details extensive analyses along four aspects. First, we compare our dense reward mechanism against the sparse PedagogicalRL reward, confirming the superiority of turn-level cognitive feedback (Appendix C.1). We then demonstrate robustness by varying backbone architectures, embedding models, and oracle choices, all yielding consistent improvements (Appendix C.2, C.3, and C.4). Next, we evaluate the sensitivity of key hyperparameters, including λ , α , δ , c , and the

number of rollouts, confirming that our defaults optimally balance cognitive scaffolding with leakage prevention (Appendix C.5, C.6, C.7, C.8, and C.9). Finally, we provide broader evaluations covering general mathematical reasoning preservation (Appendix C.10), human evaluations (Appendix C.11), qualitative case studies against GPT-4o, efficiency analysis, and additional results on MathDial (Appendix C.12, C.13, and C.14).

5 Conclusion

In this paper, we propose UCO, a cognition-oriented multi-turn reinforcement learning method for adaptive teaching. To address limitations in existing methods that lack true adaptive capability, we model the teaching process as continuous teacher-student interactions and design two synergistic reward functions. The Progress Reward quantifies cognitive progress through entropy reduction, evaluating whether students transition from uncertainty to understanding. The Scaffold Reward dynamically identifies each student’s Zone of Proximal Development and encourages teaching strategies that maintain productive struggle within this zone. Extensive experiments on BigMath and MathTutorBench benchmarks demonstrate our UCO model surpasses all equivalent-scale models and achieves performance comparable to advanced

closed-source models. We argue that modeling the teaching process from the perspective of cognitive state evolution applies not only to mathematical tutoring but also to other educational domains such as programming education and scientific reasoning. In future work, we will explore extending UCO to broader educational scenarios.

Limitations

We conducted a comprehensive evaluation of the teaching effectiveness of the UCO method, but it has the following limitations.

1. We focus on K-12 math tutoring because math problems provide clear answers and verifiable reasoning, enabling precise signal design for reinforcement learning rewards. While the teaching modeling principle based on cognitive state evolution is universally applicable, extending it to other domains faces challenges in reward design and benchmark construction. Future work should develop domain-specific cognitive progress metrics and scaffolding hierarchies for different subjects, and collaborate with researchers in education and psychology to establish evaluation benchmarks.
2. Our study employs text-based dialogue interaction, aligning with mainstream applications of large language models. However, teaching scenarios involve multimodal information such as graphics, handwritten calculations, and vocal intonation. Future work should extend the UCO framework to multimodal teaching scenarios to leverage such information.
3. The effectiveness of UCO has been validated mainly in offline datasets, and it has not yet been tested in real classrooms. The teaching impact of UCO requires further verification in real-world teaching scenarios.

Acknowledgments

This study was funded by Guangxi Science and Technology Program (2025AB25069309), the Open Research Fund of Key Laboratory of Advanced Theory and Application in Statistics and Data Science (East China Normal University). This work was supported in part by "Pioneer" and "Leading Goose" R&D Program of Zhejiang (2025C02037) and the Key R&D Program of

Hangzhou (2025SZDA0254). This work was supported in part by the National Natural Science Foundation of China (Grant No. 62506319).

References

- Alon Albalak, Duy Phung, Nathan Lile, Rafael Rafailov, Kanishk Gandhi, Louis Castricato, Anikait Singh, Chase Blagden, Violet Xiang, Dakota Mahan, et al. 2025. Big-math: A large-scale, high-quality math dataset for reinforcement learning in language models. *arXiv preprint arXiv:2502.17387*.
- Joan Bliss, Mike Askew, and Sheila Macrae. 1996. Effective teaching and learning: Scaffolding revisited. *Oxford review of Education*, 22(1):37–61.
- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation. *arXiv preprint arXiv:2402.03216*.
- Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jinheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S Yu, et al. 2025. Llm agents for education: Advances and applications. *arXiv preprint arXiv:2503.11733*.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Jason V Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S Dhillon. 2007. Information-theoretic metric learning. In *Proceedings of the 24th international conference on Machine learning*, pages 209–216.
- Vincenzo De Marco, Filippo Sciarrone, and Marco Temperini. 2024. Tutorchat: a chatbot for the support to dyslexic learner’s activity through generative ai. In *2024 IEEE International Conference on Advanced Learning Technologies (ICALT)*, pages 155–157. IEEE.
- David Dinucu-Jianu, Jakub Macina, Nico Daheim, Ido Hakimi, Iryna Gurevych, and Mrinmaya Sachan. 2025. From problem-solving to teaching problem-solving: Aligning llms with pedagogy using reinforcement learning. *arXiv preprint arXiv:2505.15607*.
- Fahmid Morshed Fahid, Jonathan Rowe, Yeojin Kim, Shashank Srivastava, and James Lester. 2024. Online reinforcement learning-based pedagogical planning for narrative-centered learning environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 23191–23199.

- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Zhi Hong, Qian Zhang, Jiahang Sun, Zhiwei Shang, Mingze Kong, Xiangyi Wang, Yao Shu, and Zhongxiang Dai. 2026. MASPOB: Bandit-based prompt optimization for multi-agent systems with graph neural networks. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, Seoul, South Korea. PMLR.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Rui Jia, Ruiyi Lan, Fengrui Liu, Zhongxiang Dai, Bo Jiang, Jing Shao, Jingyuan Chen, Guandong Xu, Fei Wu, and Min Zhang. 2026a. Castle: A comprehensive benchmark for evaluating student-tailored personalized safety in large language models. In *EMNLP*.
- Rui Jia, Yuang Wei, Ruijia Li, Yuan-Hao Jiang, Xinyu Xie, Yaomin Shen, Min Zhang, and Bo Jiang. 2026b. Diacdm: cognitive diagnosis in teacher-student dialogues using the initiation-response-evaluation framework. In *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5186–5190. IEEE.
- Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivi re, Louis Rouillard, et al. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 4.
- Martin Klissarov, R Devon Hjelm, Alexander T Toshev, and Bogdan Mazouze. On the modeling capabilities of large language models for sequential decision making. In *The Thirteenth International Conference on Learning Representations*.
- Mingze Kong, Zikun Qu, Zhongquan Zhou, Pengyu Liang, Xiang Li, Zhiwei Shang, Zhi Hong, Kaiyu Huang, Zhiyong Wang, and Zhongxiang Dai. 2026. Workflow-R1: Group sub-sequence policy optimization for multi-turn workflow construction. *arXiv preprint arXiv:2602.01202*.
- Aayush Kucheria, Nitin Sawhney, and Arto Hellas. 2025. Comparing behavioral patterns of llm and human tutors: A population-level analysis with the cima dataset. In *Workshop on Innovative Use of NLP for Building Educational Applications*. Association for Computational Linguistics.
- Wendi Li and Yixuan Li. Process reward model with q-value rankings. In *The Thirteenth International Conference on Learning Representations*.
- Yubo Li, Xiaobin Shen, Xinyu Yao, Xueying Ding, Yidi Miao, Ramayya Krishnan, and Rema Padman. 2025. Beyond single-turn: A survey on multi-turn interactions with large language models. *arXiv preprint arXiv:2504.04717*.
- Minpeng Liao, Chengxi Li, Wei Luo, Wu Jing, and Kai Fan. 2024. Mario: Math reasoning with code interpreter output-a reproducible pipeline. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 905–924.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, Bochao Wu, Bowei Zhang, Chaofan Lin, Chen Dong, et al. 2025. Deepseek-v3. 2: Pushing the frontier of open large language models. *arXiv preprint arXiv:2512.02556*.
- Hongwei Liu, Zilong Zheng, Yuxuan Qiao, Haodong Duan, Zhiwei Fei, Fengzhe Zhou, Wenwei Zhang, Songyang Zhang, Dahua Lin, and Kai Chen. 2024b. Mathbench: Evaluating the theory and application proficiency of llms with a hierarchical mathematics benchmark. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 6884–6915.
- Jiayu Liu, Zhenya Huang, Tong Xiao, Jing Sha, Jinze Wu, Qi Liu, Shijin Wang, and Enhong Chen. 2024c. Socraticlm: Exploring socratic personalized teaching with large language models. *Advances in Neural Information Processing Systems*, 37:85693–85721.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jian-Guang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, Yansong Tang, et al. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. In *The Thirteenth International Conference on Learning Representations*.
- Jakub Macina, Nico Daheim, Sankalan Chowdhury, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023a. Mathdial: A dialogue tutoring dataset with rich pedagogical properties grounded in math reasoning problems. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5602–5621.
- Jakub Macina, Nico Daheim, Ido Hakimi, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2025. Math-tutorbench: A benchmark for measuring open-ended pedagogical capabilities of llm tutors. *arXiv preprint arXiv:2502.18940*.
- Jakub Macina, Nico Daheim, Lingzhi Wang, Tanmay Sinha, Manu Kapur, Iryna Gurevych, and Mrinmaya Sachan. 2023b. Opportunities and challenges in neural dialog tutoring. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2357–2372.

- Hao Peng, Yunjia Qi, Xiaozhi Wang, Zijun Yao, Bin Xu, Lei Hou, and Juanzi Li. 2025. Agentic reward modeling: Integrating human preferences with verifiable correctness signals for reliable reward systems. *arXiv preprint arXiv:2502.19328*.
- Gabriel Poesia, David Broman, Nick Haber, and Noah Goodman. 2024. Learning formal mathematics from intrinsic motivation. *Advances in Neural Information Processing Systems*, 37:43032–43057.
- Ori Press, Ravid Shwartz-Ziv, Yann Lecun, and Matthias Bethge. 2024. The entropy enigma: Success and failure of entropy minimization. In *International Conference on Machine Learning*, pages 41064–41085. PMLR.
- Romain Puech, Jakub Macina, Julia Chatain, Mrinmaya Sachan, and Manu Kapur. 2024. Towards the pedagogical steering of large language models for tutoring: A case study with modeling productive failure. *arXiv preprint arXiv:2410.03781*.
- Zikun Qu, Min Zhang, Mingze Kong, Xiang Li, Zhiwei Shang, Zhiyong Wang, Yikun Ban, Shuang Qiu, Yao Shu, and Zhongxiang Dai. 2026. T-POP: Test-time personalization with online preference feedback. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*, Seoul, South Korea. PMLR.
- Anna Riedmann, Philipp Schaper, and Birgit Lugin. 2025. Reinforcement learning in education: A systematic literature review. *International Journal of Artificial Intelligence in Education*, pages 1–55.
- Barak Rosenshine, Carla Meister, et al. 1992. The use of scaffolds for teaching higher-level cognitive strategies. *Educational leadership*, 49(7):26–33.
- Alexander Scarlatos, Naiming Liu, Jaewook Lee, Richard Baraniuk, and Andrew Lan. 2025. Training llm-based tutors to improve student learning outcomes in dialogues. In *International Conference on Artificial Intelligence in Education*, pages 251–266. Springer.
- Karim Shabani, Mohamad Khatib, and Saman Ebadi. 2010. Vygotsky’s zone of proximal development: Instructional implications and teachers’ professional development. *English language teaching*, 3(4):237–248.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Yu Sheng, Linjing Li, and Daniel Dajun Zeng. 2025. Learning theorem rationale for improving the mathematical reasoning capability of large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 15151–15159.
- Anaïs Tack and Chris Piech. 2022. The ai teacher test: Measuring the pedagogical ability of blender and gpt-3 in educational dialogues. In *Proceedings of the 15th International Conference on Educational Data Mining*, page 522.
- LearnLM Team, Abhinav Modi, Aditya Srikanth Veerubhotla, Aliya Rysbek, Andrea Huber, Brett Wiltshire, Brian Veprek, Daniel Gillick, Daniel Kasenberg, Derek Ahmed, et al. 2024. Learnlm: Improving gemini for learning. *arXiv preprint arXiv:2412.16429*.
- Qwen Team. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2:3.
- Janneke Van de Pol, Monique Volman, and Jos Beishuizen. 2010. Scaffolding in teacher–student interaction: A decade of research. *Educational psychology review*, 22(3):271–296.
- Jiangxing Wang, Jiachen Li, Xiao Han, Deheng Ye, and Zongqing Lu. 2024a. Language model adaption for reinforcement learning with natural language action space. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1620–1634.
- Liang Wang, Nan Yang, Xiaolong Huang, Linjun Yang, Rangan Majumder, and Furu Wei. 2024b. Multilingual e5 text embeddings: A technical report. *arXiv preprint arXiv:2402.05672*.
- Rose Wang, Pawan Wirawarn, Kenny Lam, Omar Khattab, and Dorottya Demszky. 2024c. Problem-oriented segmentation and retrieval: Case study on tutoring conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 12654–12672.
- Wei Xiong, Chengshuai Shi, Jiaming Shen, Aviv Rosenberg, Zhen Qin, Daniele Calandriello, Misha Khalman, Rishabh Joshi, Bilal Piot, Mohammad Saleh, et al. Building math agents with multi-turn iterative preference learning. In *The Thirteenth International Conference on Learning Representations*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Zhangyue Yin, Qiushi Sun, Zhiyuan Zeng, Qinyuan Cheng, Xipeng Qiu, and Xuan-Jing Huang. 2025. Dynamic and generalizable process reward modeling. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4203–4233.
- Min Zhang, Bo Jiang, Jie Zhou, Yimeng Liu, and Xin Lin. 2025a. Codol: Conditional domain prompt learning for out-of-distribution generalization. In *TMLR 2026*.
- Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, et al. 2025b. Qwen3

embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*.

Zhenru Zhang, Chujie Zheng, Yangzhen Wu, Beichen Zhang, Runji Lin, Bowen Yu, Dayiheng Liu, Jingren Zhou, and Junyang Lin. 2025c. The lessons of developing process reward models in mathematical reasoning. *arXiv preprint arXiv:2501.07301*.

Li Zhong, Zilong Wang, and Jingbo Shang. 2024. Debug like a human: A large language model debugger via verifying runtime execution step by step. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 851–870.

Zihao Zhou, Bin Hu, Chenyang Zhao, Pu Zhang, and Bin Liu. 2024. Large language model as a policy teacher for training reinforcement learning agents. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 5671–5679.

A The Use of Large Language Models

In this paper, we employ LLMs as an auxiliary tool. We utilized LLMs for language polishing to enhance clarity, and we manually reviewed all modifications.

B Implementation Details

B.1 Training Configuration

We conduct all experiments on a server equipped with 8 NVIDIA H200 GPUs. The experimental environment uses Python 3.11, CUDA 12.8, and PyTorch 2.7. Training takes approximately 112 hours for the complete process.

B.2 Model Architecture

Table 5 summarizes the model configurations.

Component	Model & Purpose
Teacher	Qwen2.5-7B-Instruct for π_{teacher}
Student	Qwen2.5-14B-Instruct for $\pi_{\text{student}}^{\text{fixed}}$
Generator	Gemini-2.5-Pro for candidate sampling

Table 5: Model Architecture Configuration

B.3 Hyperparameters

All hyperparameters are determined through grid search, as shown in Table 6.

Category	Parameter	Value
GRPO Optimization	Learning rate η	5×10^{-6}
	Batch size $ \mathcal{B} $	16
	Rollouts G	4
	Max length T	10
	Discount γ	0.95
	KL weight ξ	0.01
	Norm. ϵ	10^{-8}
	Optimizer	AdamW
Progress Reward	Balance λ	0.5
	Temperature α	2.0
	Bias δ	0.7
	Embedding	BGE
Scaffold	Penalty c	0.2

Table 6: Hyperparameter Configuration

B.4 Baseline Models

Table 7 lists all baseline models in our experiments.

B.5 Evaluation Metrics

Tables 8 and 9 show the evaluation metrics.

Category	Model	Size
Closed	GPT-4o (Hurst et al., 2024)	-
	LearnLM 1.5 Pro (Team et al., 2024)	-
	LearnLM 2.0 Flash (Team et al., 2024)	-
Open	Qwen2.5-Instruct (Team, 2024)	7B
	Qwen2.5-Instruct (Team, 2024)	14B
	Qwen2.5-Instruct (Team, 2024)	72B
	DeepSeek-V3 (Liu et al., 2024a)	671B
Education	Qwen2.5-SFT (Macina et al., 2023a)	7B
	Qwen2.5-MDPO (Xiong et al.)	7B
	SocraticLM (Liu et al., 2024c)	7B
	Qwen2.5-RL (Dinucu-Jianu et al., 2025)	7B

Table 7: Baseline Model Specifications

Metric	Description
ΔSolve	Student improvement before/after dialogue
Leak Solution	LLM-detected solution divulgence
Ped-RM	Pedagogical quality scores

Table 8: BigMath Test Set Metrics

C Additional Ablation Studies

C.1 Impact of Reward Mechanisms on Training Efficiency.

Table 11 presents a comparison on the BigMath dataset to validate the superiority of our reward design. We ensure a strictly fair comparison by maintaining an identical training pipeline (same GRPO algorithm, 2,000 samples, and 4 rollouts) and only substituting our mechanism with the PedagogicalRL reward. Results show that PedagogicalRL underperforms UCO across all dimensions. Despite its explicit LLM-based hard penalties, PedagogicalRL exhibits higher Leak Solution rates, while its ΔSolve Rate and teaching quality scores (Ped-RM) lag substantially behind our method.

This performance discrepancy highlights the fundamental difference between sparse and dense reward designs. Specifically, PedagogicalRL relies on sparse, conversation-level rewards tied to the correctness of the final answer, coupled with all-or-nothing penalties for pedagogical violations. Such sparse feedback provides insufficient gradient signals, making it difficult to guide the model toward the optimal policy equilibrium without leaking answers. In contrast, UCO mitigates this by providing dense, turn-level feedback: the Progress Reward evaluates the reduction in entropy of cognitive

Domain	Task	Metric
Math Expertise	Problem Solving	Answer accuracy
	Socratic Quest.	BLEU score
Student Understand.	Sol. Correctness	F1 micro
	Mistake Location	F1 score
	Mistake Correct.	Accuracy
Pedagogy	Scaffold Gen.	Win rate
	Pedagogical IF	Win rate
	Hard versions	Win rate

Table 9: MathTutorBench Evaluation Tasks

states, and the Scaffold Reward precisely demarcates the Zone of Proximal Development (ZPD) at each step. By supplying these fine-grained, cognition-oriented signals, UCO achieves this pedagogical equilibrium, mastering adaptive teaching while strictly preventing answer leakage.

C.2 Impact of Backbone Model.

To evaluate the generalization capability of UCO across various model architectures, we conducted experiments using LLaMA3.1-8B-Instruct (Grattafiori et al., 2024) and LLaMA3.2-3B-Instruct (Grattafiori et al., 2024) as backbone models. While keeping all other experimental settings constant, we validated the performance on MathTutorBench using 500 training samples. As shown in Table 10, UCO consistently improves the performance of backbone models across different architectures and parameter scales. This transferability stems from the design logic of UCO reward functions: both *Progress Reward* and *Scaffold Reward* quantify cognitive states by analyzing the output probability distributions of the student model, rather than relying on architecture-specific features. Consequently, these feedback signals, grounded in the evolution of cognitive states, exhibit high universality, enabling UCO to generalize effectively across models of diverse scales and structures.

C.3 Impact of Embedding Models.

To verify UCO’s sensitivity to the choice of embedding models, we replace the BGE-M3 with two architecturally distinct models of comparable parameter scale: Multilingual-E5-Large-Instruct (**Multilingual-E5**) (Wang et al., 2024b) and Qwen3-Embedding-0.6B (**Qwen3-Emb-0.6B**) (Zhang et al., 2025b). As shown in Table 13 and Table 12, no single embedding model achieves comprehensive superiority across all eval-

uation dimensions. Regardless of the embedding model employed, the overall performance of UCO substantially surpasses that of the original base model. This result compellingly confirms that UCO is robust to the selection of embedding models.

C.4 Impact of Oracle Model Selection

To evaluate the sensitivity of UCO to the choice of Oracle models, we conducted comparative experiments using three distinct models: the closed-source Gemini-2.5-Pro (Comanici et al., 2025), the open-source DeepSeek-V3.2 (Liu et al., 2025), and Qwen3-235B-A22B-Instruct (**Qwen3-235B-A22B**) (Yang et al., 2025). As shown in Table 14 and Table 15, these models demonstrate highly consistent performance across core dimensions, including the improvement in solve rates, mathematical expertise, and student understanding. The variations among the models are marginal and primarily confined to the dimension of pedagogical quality. As detailed in the tables, Gemini-2.5-Pro holds a slight advantage in minimizing solution leakage (16.4%) and achieving higher overall pedagogical win rates against the backbone. This advantage is likely attributable to a richer integration of educational corpora within its pre-training data. Nevertheless, the UCO, driven by open-source models, not only substantially and comprehensively outperforms the original backbone model across all metrics but also maintains a gap relative to the closed-source frontier model that remains within an acceptable range. Overall, the experimental results verify that UCO exhibits strong robustness to the choice of Oracle models.

C.5 Impact of Rollout on Training Efficiency

We conduct comprehensive ablation studies on rollout numbers using a smaller training set of 100 samples. Figure 3 shows the impact of different rollouts on three key metrics. As rollout numbers increase progressively from 2 to 8, performance improves consistently. Figure 3(a) shows that Δ Solve Rate increases from 14.6% to 20.2%. Figure 3(b) demonstrates that Leak Solution decreases from 25.2% to 18.8%. Figure 3(c) illustrates that PedRM scores also improve steadily, with micro scores rising from 0.6 to 2.3 and macro scores from 0.5 to 2.1. However, training cost grows substantially with rollout numbers. Figure 3(d) shows the detailed training time comparison. With 100 training samples, rollout=2 takes approximately 2.9 hours. Rollout=8 requires 10.8 hours, a 3.7-fold

Model	Math Expertise		Student Understanding			Pedagogy			
	Problem solving acc.	Socratic questioning BLEU	Solution correctness F1	Mistake location micro F1	Mistake correction acc.	scaff. win rate	ped.IF win rate	scaff. [hard] win rate	ped.IF [hard] win rate
<i>LLaMA3.2-3B-Instruct as Backbone Model</i>									
LLaMA3.2-3B-Instruct	0.60	0.29	0.67	0.41	0.13	0.64	0.63	0.45	0.40
LLaMA3.2-3B-Instruct + UCO	0.63	0.30	0.69	0.44	0.16	0.69	0.68	0.51	0.46
<i>LLaMA3.1-8B-Instruct as Backbone Model</i>									
LLaMA3.1-8B-Instruct	0.70	0.29	0.63	0.29	0.09	0.61	0.67	0.46	0.49
LLaMA3.1-8B-Instruct + UCO	0.75	0.30	0.67	0.34	0.13	0.67	0.72	0.53	0.55

Table 10: Cross-architecture generalizability on MathTutorBench using 500 training samples.

Reward Function	Δ Solve Rate (%) ↑	Leak Solution (%) ↓	Ped-RM micro/macro ↑
<i>Backbone Model</i>			
Qwen2.5-7B-Instruct	11.3	29.3	-0.2/-0.5
PedagogicalRL Reward	24.3	19.2	2.8/2.7
UCO Reward (Ours)	30.2	12.9	4.6/4.5

Table 11: Comparison between PedagogicalRL and UCO reward on BigMath.

increase. We choose rollout=4 as our default setting, which achieves comparable performance to rollout=8 while reducing training time by 52% (5.2 hours vs 10.8 hours). This provides the optimal trade-off between performance and efficiency.

C.6 Impact of Balance Coefficient λ

Table 16 examines how λ balances potential capacity and semantic quality in progress rewards. At low values ($\lambda = 0.1$ and 0.3), semantic quality dominates the reward signal. Rewards depend only on output correctness. Even when the student internally understands correct reasoning, rewards stay low if the output is inaccurate. It prevents the teacher from adjusting teaching strategies to guide the student. This imbalance restricts the model’s ability to recognize internal cognitive improvements effectively. It limits improvements in Δ Solve Rate and teaching quality scores. At high values ($\lambda = 0.7$ and 0.9), potential capacity dominates the reward signal. Rewards mainly depend on increases in internal confidence and ignore output accuracy. To quickly boost student confidence, the teacher may directly provide answer-related information rather than guide autonomous reasoning. It increases answer leakage rates while decreasing Δ Solve Rate and Ped-RM scores. At $\lambda = 0.5$, performance reaches the optimum. The potential capacity component measures student confidence

in correct reasoning. The semantic quality component verifies output correctness. Together, they enable the teacher to optimize both reasoning capability and expression quality. It achieves the best performance across all metrics.

C.7 Impact of Temperature Coefficient α

Table 17 examines how the temperature coefficient α affects reward signal strength in potential ability scoring. At lower values ($\alpha = 1.0, 1.5$), the tanh function produces weak reward signals. Consequently, these signals cannot effectively distinguish confidence levels in student responses. This insufficient discrimination prevents the teacher model from identifying meaningful cognitive progress. As a result, without clear signals, the model cannot guide students toward correct reasoning paths, leading to poor teaching performance across all evaluation metrics. At higher values ($\alpha = 2.5, 3.0$), reward signals become overly sensitive to small confidence fluctuations. They amplify minor probability changes into large reward differences. This excessive sensitivity introduces noise into the training, which disrupts stability and leads to inconsistent teaching behavior. Such instability compromises the model’s overall performance. At the optimal value ($\alpha = 2.0$), the temperature coefficient achieves an ideal balance. Reward signals maintain sufficient discriminative strength while preserving stability against noise amplification. This balance enables the teacher model to reliably identify genuine cognitive improvements and deliver consistent, effective teaching strategies.

C.8 Impact of Semantic Quality Threshold δ

Table 18 examines how the bias term δ affects output filtering in semantic quality scoring. At lower values ($\delta = 0.5, 0.6$), the threshold is too lenient. Consequently, the model cannot effectively distinguish genuine understanding from shallow pattern

Model	Math Expertise		Student Understanding			Pedagogy			
	Problem solving acc.	Socratic questioning BLEU	Solution correctness F1	Mistake location micro F1	Mistake correction acc.	scaff. win rate	ped.IF win rate	scaff. [hard] win rate	ped.IF [hard] win rate
Qwen2.5-7B-Instruct	0.87	0.23	0.63	0.39	0.04	0.37	0.60	0.45	0.56
UCO w/ BGE-M3 (default)	0.89	0.24	0.65	0.40	0.06	0.60	0.66	0.63	0.68
UCO w/ multilingual-e5	0.88	0.23	0.64	0.40	0.06	0.59	0.67	0.64	0.70
UCO w/ Qwen3-Emb-0.6B	0.89	0.24	0.65	0.39	0.05	0.62	0.66	0.63	0.69

Table 12: Performance on MathTutorBench across different embedding models (500 training samples).

Model	Δ Solve Rate (%) $\uparrow$	Leak Solution (%) $\downarrow$	Ped-RM micro/macro $\uparrow$
Qwen2.5-7B-Instruct	11.3	29.3	-0.2 / -0.5
UCO w/ BGE-M3 (default)	22.8	16.4	3.6 / 3.5
UCO w/ Multilingual-E5	22.5	16.7	3.7 / 3.7
UCO w/ Qwen3-Emb-0.6B	23.2	16.9	3.5 / 3.4

Table 13: Performance on BigMath across different embedding models (500 training samples).

Model	Δ Solve Rate (%) $\uparrow$	Leak Solution (%) $\downarrow$	Ped-RM micro/macro $\uparrow$
Qwen2.5-7B-Instruct	11.3	29.3	-0.2 / -0.5
UCO w/ Gemini-2.5-Pro	22.8	16.4	3.6 / 3.5
UCO w/ DeepSeek-V3.2	23.1	17.3	3.2 / 3.1
UCO w/ Qwen3-235B-A22B	22.6	17.6	3.0 / 2.9

Table 14: Comprehensive evaluation on BigMath across various oracle models (500 training samples).

matching. This causes the reward function to inappropriately reward superficial matches lacking true reasoning, resulting in teaching strategies that lack precision and effectiveness. At higher values ($\delta = 0.8, 0.9$), the threshold becomes overly strict. Consequently, many responses demonstrating valid understanding are incorrectly filtered out. This causes the reward signal to inadequately capture the actual progression in students’ cognitive development. As a result, the teacher model tends to provide excessive scaffolding, which undermines students’ autonomous reasoning capacity and increases answer leakage risk. At the optimal value $\delta = 0.7$, the threshold achieves an ideal balance. Unlike lower values that allow noisy signals to mislead the optimization process, or higher values that discard valuable learning signals, this optimal setting enables the model to accurately filter ambiguous outputs while reliably identifying genuine understanding. Consequently, the reward function precisely evaluates students’ cognitive states, and performance is optimized across all metrics.

C.9 Impact of Deviation Coefficient c .

Table 19 examines how the deviation penalty coefficient c regulates scaffolding precision. This coefficient penalizes teaching strategies that deviate from students’ Zone of Proximal Development (ZPD). At lower values ($c = 0.1$), the penalty strength is insufficient. The model receives weak feedback signals for scaffolding deviations. The teacher model struggles to effectively adjust teaching difficulty to appropriate levels. All evaluation metrics show suboptimal performance. At higher values ($c = 0.3, 0.4, 0.5$), penalty strength becomes overly strict. The model’s teaching strategies are over-constrained. The teacher model’s ability to respond to individual student needs declines accordingly. Limited teaching flexibility leads to overall performance degradation. At the optimal value ($c = 0.2$), penalty strength achieves an ideal balance. The model receives sufficient incentive to maintain scaffolding precisely within the ZPD range. All evaluation metrics achieve optimal performance under this configuration.




Figure 3: Impact of rollout numbers on model performance and training efficiency.

C.10 General Mathematical Reasoning Ability Preservation

Table 20 shows UCO’s performance on general math reasoning benchmarks. UCO achieves 85.9% and 75.0% accuracy on GSM8K and MATH500.

Model	Math Expertise		Student Understanding			Pedagogy			
	Problem solving acc.	Socratic questioning BLEU	Solution correctness F1	Mistake location micro F1	Mistake correction acc.	scaff. win rate	ped.IF win rate	scaff. [hard] win rate	ped.IF [hard] win rate
Qwen2.5-7B-Instruct	0.87	0.23	<i>Backbone Model</i>			0.37	0.60	0.45	0.56
UCO w/ Gemini-2.5-Pro	0.89	0.24	0.65	0.40	0.06	0.60	0.66	0.63	0.68
UCO w/ DeepSeek-V3.2	0.89	0.24	0.65	0.39	0.06	0.56	0.62	0.57	0.62
UCO w/ Qwen3-235B-A22B	0.87	0.23	0.64	0.39	0.05	0.54	0.61	0.53	0.58

Table 15: Comprehensive evaluation on MathTutorBench across various oracle models (500 training samples).

λ	Δ Solve Rate (%) ↑	Leak Solution (%) ↓	Ped-RM micro/macro ↑
0.1	26.5	16.8	3.5/3.6
0.3	27.8	15.3	3.8/4.0
0.5	30.2	12.9	4.6/4.5
0.7	29.1	13.5	4.3/4.2
0.9	27.4	14.9	3.9/3.7

Table 16: Ablation study on λ balancing potential capacity and semantic quality in progress reward.

α	Δ Solve Rate (%) ↑	Leak Solution (%) ↓	Ped-RM micro/macro ↑
1.0	27.4	15.1	3.8/4.0
1.5	28.9	13.8	4.1/4.2
2.0	30.2	12.9	4.5/4.4
2.5	28.6	14.3	4.2/4.0
3.0	27.1	16.1	3.3/2.8

Table 17: Ablation study on temperature α controlling reward sensitivity in potential capacity score.

δ	Δ Solve Rate (%) ↑	Leak Solution (%) ↓	Ped-RM micro/macro ↑
0.5	27.2	16.1	3.7/3.9
0.6	28.6	14.5	4.0/4.2
0.7	30.2	12.9	4.6/4.5
0.8	29.3	13.7	4.2/4.3
0.9	26.8	17.3	3.6/3.8

Table 18: Ablation study on semantic threshold δ for filtering ambiguous outputs in semantic quality score.

c	Δ Solve Rate (%) ↑	Leak Solution (%) ↓	Ped-RM micro/macro ↑
0.1	27.6	15.8	3.8/3.9
0.2	30.2	12.9	4.6/4.5
0.3	28.7	14.1	4.2/4.1
0.4	27.3	15.6	3.9/4.0
0.5	26.1	17.2	3.5/3.7

Table 19: Ablation study on deviation coefficient c regulating scaffolding precision in scaffolding reward.

Model	GSM8K	MATH500
<i>Backbone Model</i>		
Qwen2.5-7B-Instruct (Team, 2024)	86.8	75.4
<i>Education-specific Models</i>		
Qwen2.5-7B-SFT (Macina et al., 2023a)	79.5	66.0
Qwen2.5-7B-MDPO (Xiong et al.)	87.0	76.4
Qwen2.5-7B-RL (Dinucu-Jianu et al., 2025)	86.1	73.6
UCO	85.9	75.0

Table 20: Evaluation of general mathematical reasoning benchmarks GSM8K and MATH500.

This represents only 0.9 and 0.4 percentage point drops from the backbone model. Other education-specific models show larger performance declines. Qwen2.5-7B-SFT drops by 7.3 and 9.4 percentage points. Qwen2.5-7B-RL drops by 0.7 and 1.8 percentage points. Qwen2.5-7B-MDPO performs best on general ability. However, it shows very limited improvement in teaching quality. This demonstrates that UCO achieves the best trade-off between teaching specialization and general ability.

C.11 Human Evaluation.

To further evaluate the effectiveness of UCO, we invited three experienced teachers to conduct a human assessment. Specifically, we randomly sampled 100 problems from the BigMath test set. Following the experimental setup of PedagogicalRL (Dinucu-Jianu et al., 2025), various models acted as teachers to interact with a unified student model (Meta-Llama-3.1-8B-Instruct) to generate complete multi-turn teacher-student dialogues. The three teachers assessed the multi-turn teacher-student dialogues generated by different models against those produced by GPT-4o to determine the Win Rate. This evaluation is from five dimensions proposed by SocraticLM (Liu et al., 2024c), namely Overall Quality, Incorrect Answer Recognition Accuracy (IARA), Correct Answer Recognition Accuracy (CARA), Successful Explanation Rate (SER), and Successful Rejection Rate (SRR).

As shown in Table 21, UCO achieved a win rate of 0.69 against GPT-4o, significantly outperforming all baseline methods. Results demonstrate that UCO can accurately perceive student states and dynamically adjust teaching strategies. Compared to GPT-4o, which tends to provide solutions directly in the later stages of a conversation, UCO exhibits superior scaffolding and guidance capabilities.

Teacher Model	Win Rate (vs. GPT-4o)
<i>Backbone Model</i>	
Qwen2.5-7B-Instruct	0.24
GPT-4o	0.50
SocraticLM	0.57
Qwen2.5-7B-RL	0.61
UCO (Ours)	0.69

Table 21: Human evaluation results on BigMath.

C.12 Case Study

To verify the effectiveness of our proposed method, we conduct a detailed qualitative case study. We compare our model with GPT-4o-2024-11-20 (Hurst et al., 2024), both serving as teacher models. The student model uses Meta-Llama-3.1-8B-Instruct. We carefully select two representative problems from the BigMath (Albalak et al., 2025) test set. The maximum dialogue turns are set to 16. The teacher model outputs `<end_of_conversation>` when it determines the conversation should end. Following prior research (Dinucu-Jianu et al., 2025), the specific prompt designs for teacher and student models are shown in Table 24. We present two dialogue cases and conduct an in-depth comprehensive analysis to reveal the core advantages of our method:

For case one, this problem involves basic rate calculation. Table 22 shows the complete dialogue process of both models.

Our method demonstrates stronger adaptive guidance ability. First, our model activates students’ metacognitive thinking through open-ended questions (“What two core pieces of information do we need?”). This heuristic questioning forces students to identify the problem’s essential structure. Students do not passively accept solution steps. In contrast, GPT-4o directly states the solution direction (“we look at two cup types separately”). It deprives students of independent thinking opportunities. Second, our method exhibits more precise progressive scaffolding. When the student proposes a complete plan, the teacher model only gives affir-

UCO	GPT-4o
<i>Problem: Gina can paint 6 rose cups/hour and 7 lily cups/hour. Her store receives an order for 6 rose cups and 14 lily cups. If she gets \$90 total, what is her hourly pay?</i>	
T: To calculate a "dollars per hour" rate, what two core pieces of information do we need?	T: Let's solve this. What do you think we should calculate first?
S: The total dollars she made... and the total hours she worked?	S: Um... maybe how long it takes to finish the order?
T: Exactly. Total pay is given. So the real task is finding total hours. The order has two parts. How would you approach that?	T: Good thinking. To get total time, we look at two cup types separately. For 6 rose cups at 6/hour, that takes 1 hour. Clear?
S: I'd figure out time for roses, then lilies, and add them.	S: Clear, that's simple.
T: That's a solid plan. Go ahead, what's the time for roses?	T: The other part is 14 lily cups. She paints 7/hour. Can you calculate how long that takes?
S: She paints 6 at 6/hour... so 1 hour.	S: 14 divided by 7 is 2 hours.
T: Right. And the lilies?	T: Right. So total time is 1 + 2 = 3 hours. Now we know total time and pay, what's the final step?
S: 14 cups at 7/hour... $14 \div 7 = 2$ hours.	S: Um... use \$90 divided by 3 hours?
T: Correct. Now you have everything. What's the final step?	T: Yes, that gives her hourly rate.
S: Total time is 3 hours. So $\$90 \div 3 = \$30/\text{hour}$.	S: The result is \$30/hour.
T: You've walked through the solution logically. <code><end_of_conversation></code>	T: Right, the final answer is \$30. <code><end_of_conversation></code>

Table 22: Case One Dialogue Comparison

mation and encourages execution (“That’s a solid plan. Go ahead”). It transfers complete computational agency to the student. GPT-4o directly provides intermediate calculation results (“that takes 1 hour”). It interferes with students’ autonomous problem-solving process. Such over-guidance diminishes students’ cognitive engagement and independent reasoning. Finally, our method emphasizes students’ thinking process when ending the dialogue (“You’ve walked through the solution logically”). It strengthens students’ metacognitive reflection on solution strategies. GPT-4o only confirms answer correctness. It ignores the summary of the thinking process.

For case two, this problem examines students’ comprehensive understanding of periodicity and parity. Table 23 shows strategies of both models when handling abstract mathematical concepts.

When handling complex abstract problems, our method shows finer cognitive scaffolding construction ability. First, our model verifies students’ basic concept understanding (“Recall its definition?”). It guides students to independently derive periodicity (“What about increasing by 2?”). It promotes students to abstract derivation results into mathematical properties (“What does this property tell us?”). This structured guidance cultivates students’ mathematical reasoning ability and concept transfer abil-

UCO	GPT-4o
<i>Problem: $f(x)$ on $\mathbb{R}$ is even, satisfies $f(x+1) = -f(x)$. When $x \in [-1, 0)$, $f(x) = (\frac{1}{2})^x$. Find $f(\log_2 8)$.</i>	
T: This gives three pieces of info about $f(x)$. First, it's "even." Recall its definition? S: I think it's $f(-x) = f(x)$. T: Correct. Second, $f(x+1) = -f(x)$. This tells what happens when input increases by 1. What about increasing by 2? S: Hmm... $f(x+2) = -f(x+1)$. Since $f(x+1) = -f(x)$, then $f(x+2) = f(x)$. T: Your reasoning is correct. What does this property tell us? S: The function is periodic with period 2. T: Very good. Now what we need: $f(\log_2 8)$. What is $\log_2 8$? S: 2 to power 3 is 8, so it's 3. We need $f(3)$. T: Right. How can we use the period property? S: $f(3) = f(1)$. T: Good. Now the problem becomes finding $f(1)$. Which condition helps? S: $f(x+1) = -f(x)$? T: Yes, how would you use it? S: If I let $x = 0$, I get $f(1) = -f(0)$. T: Excellent. How can we find $f(0)$? S: If I let $x = -1$, I get $f(0) = -f(-1)$. T: Perfect! What is $f(-1)$? S: $f(-1) = (1/2)^{-1} = 2$. T: Correct. So what is $f(0)$? S: $f(0) = -2$. T: Can you now connect all results? S: Since $f(0) = -2$, then $f(1) = 2$. And $f(3) = 2$. T: Exactly right! You successfully applied multiple function properties to complete the derivation. <end_of_conversation>	T: This looks complex, but the key is using given properties. We need $f(\log_2 8)$, which is $f(3)$. S: But the formula's domain is $[-1, 0)$, and 3 isn't in it. T: Correct. We use properties to transform $f(3)$. Which helps? S: I'm not sure... the periodic one? T: Yes. We deduce $f(x+2) = f(x)$, so period is 2. Using this, $f(3) = f(1)$. S: Now we need $f(1)$. T: Let's use $f(x+1) = -f(x)$. When $x = 0$, this becomes $f(1) = -f(0)$. S: Oh, so we just need $f(0)$. T: Right. We use $f(x+1) = -f(x)$ again. When $x = -1$, we get $f(0) = -f(-1)$. S: So $f(-1) = (1/2)^{-1} = 2$. T: Correct. So $f(0) = -2$. S: Then $f(1) = -f(0) = 2$. T: Exactly. Since $f(3) = f(1)$, what's the final answer? S: The final answer is 2. T: Correct. The answer is 2. <end_of_conversation>

Table 23: Case Two Dialogue Comparison

ity. GPT-4o skips the concept verification stage. It directly tells students “We deduce $f(x+2) = f(x)$ ” and “period is 2”. While expedient, this approach limits students’ opportunities for independent mathematical discovery. It hinders the formation of a deep understanding. Second, our method requires active student participation at each reasoning node (“Which condition helps?” “How would you use it?”). This continuous cognitive activation ensures students remain in the leading position of problem-solving. GPT-4o tends to provide explicit procedural directions (“Let’s use $f(x+1) = -f(x)$ ”). Students only need to execute the given steps. Finally, our method strengthens students’ overall understanding of multi-step reasoning through comprehensive questions (“Can you now connect all

Teacher Prompt

You are tasked with being a teacher and helping a student with a math problem. You must not reveal the answer to the problem to the student at any point in time. Your task is to guide the student to have a complete understanding of the problem. Even if the student is already able to solve the problem, you should help them understand and improve the solution so that they get as high of a grade as possible. If possible, do not respond with overly long responses to the student. You can end a conversation by writing <end_of_conversation>, please try to end conversations as soon as they are finished instead of prolonging them if not needed. But do not end them prematurely either. Here is the math problem: {{ problem }}.

Student Prompt

You will act as a student in a conversation with a teacher in training. You will need to act as much like a student as possible. If possible do not respond with overly long messages. The conversation with the teacher will be about this math problem: {{ problem }}. You may or may not know how to solve it already, let the teacher guide you to the correct understanding. You will be tested at the end and scored thus it is best if you collaborate with the teacher as it has more experience in math than you.

Table 24: Teacher and Student Prompts

results?”). It facilitates the transfer of the application to similar problems.

Through comparative analysis of these two cases, our method demonstrates stronger adaptive teaching ability. It reflects the effectiveness of the proposed reward function.

C.13 Efficiency Analysis.

We analyze the efficiency of UCO from three dimensions: training cost, inference latency, and scalability. During training, UCO requires approximately 112 hours on 8 NVIDIA H200 GPUs to complete the full training process, compared to 48 hours for Qwen2.5-7B-RL, with the additional overhead primarily stemming from interactions with the oracle model during each dialogue turn. However, this cost is incurred during the training phase. At inference time, UCO is essentially a standard Qwen2.5-7B-Instruct model that no longer relies on the oracle model or the student model, so its inference latency is fully consistent with baselines of the same scale, with single-turn teacher response generation taking approximately 0.3–0.5 seconds on a single H200 GPU, meeting the interaction requirements of real-time tutoring. Furthermore, the 7B parameter size allows UCO to perform efficient inference on a single GPU, enabling it to provide tutoring services at low hardware cost.

C.14 Evaluation on Additional Benchmarks

To validate the effectiveness of UCO across broader evaluation benchmarks, we adopted the dual-dimensional framework from Tutorbot-DPO (Scarlatos et al., 2025) on the MathDial test set, assessing models on *Student Outcomes* and *Pedagogical Principles*. As shown in Table 25, UCO surpasses Tutorbot-DPO in the overall pedagogical score, with significant advantages in advancing progress, hinting, and encouraging students. Conversely, it scores slightly lower on student outcomes and direct error correction. This performance difference reflects a fundamental divergence in teaching strategies: Tutorbot-DPO optimizes for single-turn correctness, causing the model to directly point out errors and provide explicit answers. In contrast, UCO leverages online reinforcement learning to provide progressive hints. Although the multi-turn elicitation process slightly decreases the immediate correctness metric, it enables the model to dynamically adjust its scaffolding based on the student’s state. This deliberate trade-off, which prioritizes long-term cognitive engagement over short-term immediate metrics, better aligns with established educational principles.

Method	Student	Pedagogical Principles						Overall
	Outcomes	Acc.	Prog.	Err.	Hint	Wth.	Enc.	Score
	↑	↑	↑	↑	↑	↑	↑	↑
Tutorbot-DPO	0.65	0.97	0.96	0.92	0.96	0.99	0.92	9.37
UCO	0.63	0.95	0.98	0.89	0.99	0.99	0.98	9.48

Table 25: Evaluation of student outcomes and pedagogical principles on the MathDial benchmark. Pedagogical metrics include Accuracy (**Acc.**), Progress (**Prog.**), Error identification (**Err.**), Strategic Hinting (**Hint**), Withholding the answer (**Wth.**), and Encouraging (**Enc.**).