

DPRM: A Dual Implicit Process Reward Model in Multi-Hop Question Answering

Xinyi Wang¹, Yiping Song¹, Zhiliang Tian^{1*}, Bo Liu^{2*}, Tingjin Luo¹, Minlie Huang³

¹National University of Defense Technology

²Academy of Military Sciences

³The CoAI Group, Department of Computer Science and Technology, Tsinghua University
tianzhiliang@nudt.edu.cn, kyle.liu@nudt.edu.cn

Abstract

In multi-hop question answering (MHQA) tasks, Chain of Thought (CoT) improves the quality of generation by guiding large language models (LLMs) through multi-step reasoning, and Knowledge Graphs (KGs) reduce hallucinations via semantic matching. Outcome Reward Models (ORMs) provide feedback after generating the final answers but fail to evaluate the process for multi-step reasoning. Traditional Process Reward Models (PRMs) evaluate the reasoning process but require costly human annotations or rollout generation. While implicit PRM is trained only with outcome signals and derives step rewards through reward parameterization without explicit annotations, it is more suitable for multi-step reasoning in MHQA tasks. However, existing implicit PRM has only been explored for plain text scenarios. When adapting to MHQA tasks, it cannot handle the graph structure constraints in KGs and capture the potential inconsistency between CoT and KG paths. To address these limitations, we propose the DPRM (Dual Implicit Process Reward Model). It trains two implicit PRMs for CoT and KG reasoning in MHQA tasks. Both PRMs, namely KG-PRM and CoT-PRM, derive step-level rewards from outcome signals via reward parameterization without additional explicit annotations. Among them, KG-PRM uses preference pairs to learn structural constraints from KGs. DPRM further introduces a consistency constraint between CoT and KG reasoning steps, making the two PRMs mutually verify and collaboratively optimize the reasoning paths. We also provide a theoretical demonstration of the derivation of process rewards. Experimental results show that our method outperforms 13 baselines on multiple datasets with up to 16.6% improvement on Hit@1.

Code — <https://github.com/suyun417/DPRM>

Extended version — <https://arxiv.org/abs/2511.08364>

1 Introduction

In the Multi-hop Question Answering (MHQA) task, Chain of Thought (CoT) helps large language models (LLMs) to generate answers by thinking step by step, and structured knowledge (Liu et al. 2025; Li et al. 2025) like Knowledge Graphs (KGs) makes reasoning more accurate by matching facts. These KG-enhanced LLM reasoning methods achieve

good results in many domains, including customer service (Xu et al. 2024), finance (Nishat et al. 2025), medical inquiry (Zhao et al. 2024; Wu et al. 2025), and other scenarios.

Most of the existing MHQA methods based on LLMs and KGs follow the paradigm of Graph Retrieval Augmented Generation (GraphRAG) (Mavromatis and Karypis 2024; He et al. 2025; Edge et al. 2024; Hu et al. 2025; Wang et al. 2025). They retrieve query-related knowledge from KG to assist LLM in generating reasoning paths and answers. Retrieval is usually completed before reasoning, and the retrieved content is fixed. Once the retrieval results are incomplete or contain redundant information, LLM has difficulty acquiring new knowledge for step-level error correction. Other works make the LLMs interact with the KGs to generate answers step by step (Sanmartin 2024; Li et al. 2024; Jin et al. 2024; Dong et al. 2024b). In interactive systems, the planning and reflection processes rely on the LLM’s ability to evaluate reasoning paths, but this evaluation is typically qualitative rather than providing quantitative scores. Additionally, KG retrievers optimize for shortest paths but lack path quality evaluation. As a result, error paths will lead to wrong answers. Therefore, a quantitative evaluation mechanism is needed to evaluate and filter CoT steps and their corresponding KG triples during reasoning.

Meanwhile, in mathematical reasoning tasks (Ziyi et al. 2025), there is a widely used and proven effective quantitative evaluation mechanism for reasoning—reward models (RMs). Early RMs are mostly explicit. The outcome reward models (ORMs) assign rewards to the entire response and provide feedback after generating the final answer (Christiano et al. 2017; Liu et al. 2024; Ouyang et al. 2022). It can effectively replace manual reward design and make the model learn human preferences. However, ORM fails to evaluate the reasoning process for complex tasks requiring multi-step reasoning. In addition, since rewards are emitted only at the terminal state, the sparse supervision signals struggle with stability and efficiency in training (Cao et al. 2024; Chan et al. 2024). Process reward models (PRMs) break through this limitation by providing evaluation of intermediate steps (Lightman et al. 2023; Wang et al. 2023b; Zhao et al. 2025; Luo et al. 2024b; Zhang et al. 2024). This fine-grained feedback improves training efficiency and final performance. However, obtaining high-quality process-level supervision signals requires high-cost manual annota-

*Corresponding authors

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

tion (Lightman et al. 2023) or demands heavy computation to generate rollouts for estimation (Wang et al. 2023b).

Implicit PRM does not rely on explicit annotation (Yuan et al. 2024; Cui et al. 2025; Guan et al. 2025). It is trained only with outcome signals and derives step rewards through reward parameterization without manual or automated annotations. Such implicit PRM is suitable for multi-step reasoning tasks such as MHQA because it can quantitatively evaluate the reasoning steps without explicit annotations. However, existing implicit PRM only focuses on the evaluation of intermediate steps in plain text scenarios (CoT), lacking consideration of the graph structure constraint of structured knowledge (KG). This overlook is critical because KGs provide explicit relational constraints to enhance LLM reasoning. Further, designing separate implicit PRMs for CoT and KG and applying them independently in reasoning may cause semantic and logical inconsistency between top-scoring paths. This inconsistency may lead to conflict when integrating different paths to get the final answer.

To solve these problems, we propose DPRM, a dual implicit process reward model for automatic annotation and collaborative reasoning. It trains two PRMs, CoT-PRM and KG-PRM, and both of them derive process rewards from outcome signals via reward parameterization without explicit annotations. This mechanism compares reasoning paths in pairs, ranking them based on whether they follow graph rules and maintain logical consistency. In training, CoT preference pairs transformed from KG data enable CoT-PRM to learn logic structure, and KG preference pairs derived from CoT data introduce relational paths absent in the original KGs. In reasoning, both PRMs act as evaluators to screen KG paths and CoTs. Based on their rewards, high-quality KG triples and CoT steps reinforce each other. This iterative interaction mechanism effectively fuses the advantages of bimodal information to make the model generate more accurate answers. In summary, our contributions are:

- We propose a framework for dual-modality PRMs in MHQA. It automatically derives CoT-PRM and KG-PRM from ORMs via mathematical derivation without manual step-level annotations. The KG-PRM learns to capture topological constraints via path preference pairs.
- We design a bidirectional KG-CoT enhancement mechanism. In training, KG-PRM and CoT-PRM enhance each other through co-training. In reasoning, the two PRMs iteratively guide KG path acquisition and CoT generation to form a self-correcting reasoning.
- Our method achieves the best on multiple datasets with up to 16.6% improvement on Hit@1.

2 Related Works

2.1 KG-enhanced LLM Reasoning

Most of the existing MHQA methods based on LLMs and KGs follow the paradigm of GraphRAG. GNN-RAG (Mavromatis and Karypis 2024) first retrieves candidate answers on the KG and then uses Graph Neural Networks (GNNs) to extract the reasoning paths. G-Retraver (He et al. 2025) first obtains relevant nodes and edges and then

finds the most relevant KG subgraph through the Rewarded Steiner Tree Algorithm. SubgraphRAG (Li, Miao, and Li 2025) generates answers by retrieving relevant subgraphs. GRAG (Hu et al. 2025) introduces k-hop self-graph matching to obtain the most relevant KG subgraph. Retrieval is usually completed before LLM reasoning with fixed content. Once the retrieval results are incomplete or redundant, LLM cannot dynamically update knowledge to supplement and correct the reasoning process.

Other works make the LLMs interact with the KG to generate answers step by step. KG-RAG (Sanmartin 2024) searches the KG following the exploration plan from the LLM. GraphReader (Li et al. 2024) uses an LLM to explore the KG, continuously updating a notebook to record relevant information. Graph-CoT (Jin et al. 2024) makes LLMs traverse the KG step by step to find the key information. RoG (Luo et al. 2024a) proposes a planning-search-reasoning framework. In interactive systems, the planning and reflection parts depend on the LLM’s evaluation, but this evaluation is qualitative rather than providing quantitative scores. Additionally, KG retrievers optimize for the shortest query-answer paths but lack path quality evaluation. If paths have errors, they will lead to wrong answers.

2.2 Reward Model

Early RMs mostly use explicit rewards. The ORMs assign rewards to the entire response and provide feedback after generating the final answer. Christiano et al. first proposes an RM framework based on human preferences, laying the foundation for ORM. Ouyang et al. proposes a three-stage RLHF framework for ORMs. Skywork-reward (Liu et al. 2024) trains the RM with high-quality small samples through efficient preference data screening. These methods can effectively replace manual reward design and make the model learn human preferences. However, ORMs fail to evaluate the reasoning process for complex tasks requiring multi-step reasoning. In addition, since rewards are emitted only at the terminal state, such sparse supervision signals struggle with stability and efficiency in training.

PRMs break through this limitation by providing evaluations of intermediate steps. Lightman et al. demonstrates that PRM solves the problem of incorrect step localization invalidation caused by sparse rewards. However, this method relies on humans to annotate reasoning steps, which is costly and difficult to generalize. Subsequent works adopt the methods of automatic annotation. Math-Shepherd (Wang et al. 2023b) first proposes an automatic annotation framework for Monte Carlo Tree Search (MCTS). It uses LLM to generate the subsequent rollouts and calculates the probability of each step leading to the correct answer. GenPRM (Zhao et al. 2025) adds code validation to this framework to filter logic error samples. OmegaPRM (Luo et al. 2024b) adds binary search to locate the first error step quickly. ReST-MCTS* (Zhang et al. 2024) uses MCTS to obtain positive samples to train PRM and then uses the trained PRM to guide the MCTS search. This fine-grained feedback improves both training efficiency and final performance. However, obtaining high-quality process-level supervision signals requires high-cost manual annotation or demands heavy

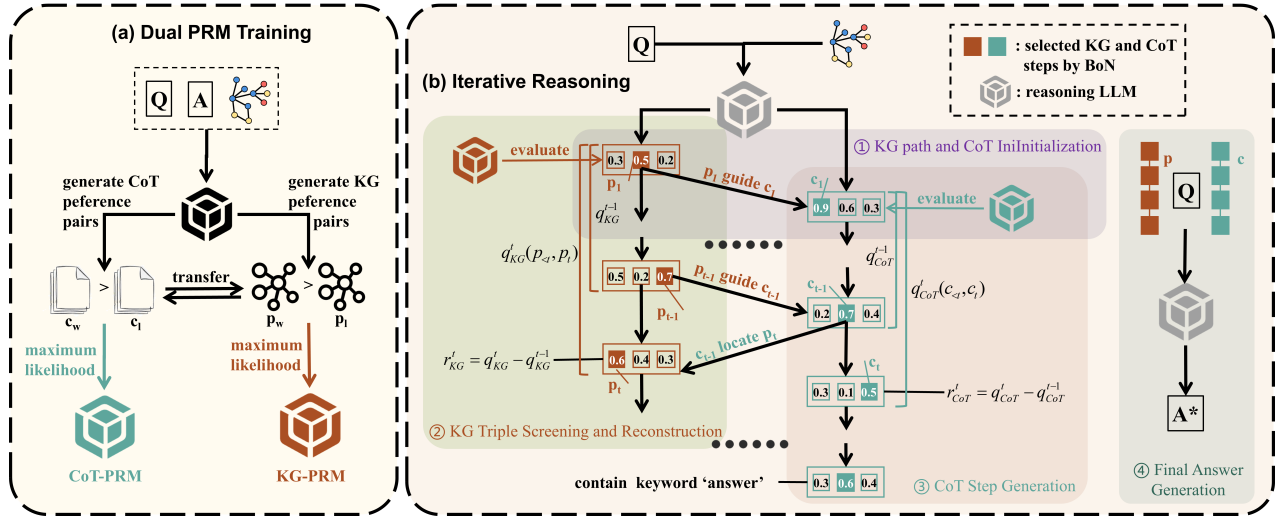


Figure 1: The overview of DPRM. (a) **Dual PRM Training** trains CoT-PRM and KG-PRM with outcome signals. KG-PRM uses preference pairs to learn structural constraints. Co-training of both PRMs makes them mutually verify and collaboratively optimize the KG paths and CoTs. (b) **Iterative Reasoning** contains 4 parts: ① KG Path and CoT Initialization, ② KG triple screening and reconstruction (use KG-PRM), ③ CoT step generation (use CoT-PRM), and ④ Final Answer Generation.

computation to generate rollouts for estimation.

Implicit PRM does not rely on explicit annotation. Implicit rewards originate from Direct Preference Optimization (DPO) (Rafailov et al. 2023). It transforms the reinforcement learning (RL) objective into a binary cross-entropy loss. Subsequent works explore how to derive PRMs through implicit rewards. ImplicitPRM (Yuan et al. 2024) is trained only with outcome signals and derives step rewards from full response rewards through reward parameterization. PRIME (Cui et al. 2025) combines implicit PRM with RL to compute step rewards and avoids reward hacking through an online mechanism. rStar-Math (Guan et al. 2025) proposes a self-evolving framework based on implicit PRM. These methods can effectively reduce the annotation cost of rewards. However, existing implicit PRM only focuses on the evaluation of intermediate steps in plain text scenarios (CoT), lacking consideration of the graph structure constraint of structured knowledge (KG). Further, it cannot cope with the collaborative reasoning process of KGs and CoTs.

3 Method

3.1 Dual PRM Architecture

Given a question Q and a related KG G , we aim to obtain an accurate answer A^* . First, we train two implicit PRMs with outcome supervision to obtain evaluators that can quantify the quality of KG paths and CoTs, called KG-PRM and CoT-PRM. Then, in reasoning, we use the two PRMs to evaluate the KG path steps and CoT steps to screen the better ones. The selected KG path steps will guide the generation of the CoT steps, and the entities in the CoT steps will help to capture the next KG path steps. They complement each other and interact iteratively to get the accurate final answers. Our method contains two key parts: **Dual PRM Training** and

Iterative Reasoning. The overall framework is in Figure 1.

Dual PRM Training In MHQA tasks, since the quality of CoT and KG reasoning is highly correlated, we adopt a data interaction strategy to make CoT-PRM and KG-PRM learn from two types of preference pairs. These PRMs use special reward functions in training so they can be trained with outcome signals and used to evaluate the reasoning process, which is theoretically demonstrated. Specifically, we first initialize CoT-PRM and KG-PRM with the data of the corresponding modality. After that, we transform KG data into CoTs to train CoT-PRM, and transform CoT data into KG paths to train KG-PRM. KG-PRM uses preference pairs to learn structural constraints from KGs.

Iterative Reasoning After training, we use the PRMs to evaluate KG paths and CoTs. In each iteration, the LLM retrieves and reconstructs new KG triples based on the question Q and the previous CoT $\{c_1, \dots, c_{i-1}\}$. The KG-PRM evaluates the N triples and selects the best as p_i . p_i will then guide the generation of the next CoT steps. The CoT-PRM evaluates the samples to select the best as c_i . The iteration continues until c_i reaches the final answers.

3.2 Dual PRM Training Framework

PRM Initialization We first use the KG and CoT data to initialize the corresponding KG-PRM and CoT-PRM. The goal of initialization is to establish foundational evaluation criteria for the two PRMs. We directly optimize the model’s preference probability distribution over preference pairs. KG-PRM learns to capture topological constraints for long chain reasoning via path preference pairs.

For KG-PRM, each piece of training data includes the question Q , true KG path p_w (w stands for win), and false KG path p_l (l stands for lose), which is (Q, p_w, p_l) . KG path

is an alternate sequence of entities and relations, starting with the question entity and ending with the answer entity.

$$\{e_Q \xrightarrow{r_1} e_1 \xrightarrow{r_2} \dots \xrightarrow{r_n} e_A\} \quad (1)$$

where e_i are entities in intermediate steps. Each step in the KG path is a triple like (e_Q, r_1, e_1) . Given the paths associated with the graph structure, the KG-PRM can learn to capture topological constraints. The true samples in the data are KG paths from the question entities to the ground truth entities. And as shown in Figure 2, the false samples are obtained in the following ways: (1) Factual errors (The entities in triples are wrong.); (2) Logical errors (The relations in triples are wrong.); (3) Logical break (The entities in adjacent triples of a KG path cannot be aligned or connected.).




Figure 2: True samples and false samples.

For a question Q , a KG path p , a KG policy model π_ϕ^{KG} , and a reference model π_{ref} , the reward function is:

$$r_{KG}(p | Q) = \gamma \log \frac{\pi_\phi^{KG}(p | Q)}{\pi_{ref}(p | Q)} \quad (2)$$

where γ is a hyperparameter used to control the strength of the preference signal. In training, the policy model is trained to be consistent with human preferences, and the reference model guarantees stability during optimization. Specifically, the loss function L_{KG} is to maximize the log probability of the reward score difference between true and false samples. For data (Q, p_w, p_l) , the loss function is:

$$L_{KG} = -\log \sigma \left(\gamma \log \frac{\pi_\phi^{KG}(p_w | Q)}{\pi_{ref}(p_w | Q)} - \gamma \log \frac{\pi_\phi^{KG}(p_l | Q)}{\pi_{ref}(p_l | Q)} \right) \quad (3)$$

where σ is the sigmoid function. Then KG-PRM can learn to identify the accurate KG paths that follow KG logic rules.

Even though KG-PRM is trained with only outcome signals, it can be used to evaluate the KG path steps because the following proposition:

Proposition 3.1. Consider an ORM where the reward is $r_\theta(y) = \beta \log \frac{\pi_\theta(y)}{\pi_{ref}(y)}$. Define $q_\theta^t(y_{<t}, y_t) := \sum_{i=1}^t \beta \log \frac{\pi_\theta(y_i | y_{<i})}{\pi_{ref}(y_i | y_{<i})}$. q_θ^t is the exponential average of r_θ at step t .

$$q_\theta^t(y_{<t}, y_t) = \beta \log \mathbb{E}_{\pi_{ref}(y | y_{\leq t})} e^{\frac{1}{\beta} r_\theta(y)}$$

Hence, q_θ^t represents an exact expectation of outcome reward r_θ at step t .

So the KG-PRM can learn the expectation of outcome reward r_{KG} at step t to calculate the process reward. Following the previous works (Lu et al. 2024), we define the process reward at step t as the difference between q^t and q^{t-1} . The process reward at step t for KG-PRM is,

$$r_{KG}^t = q_{KG}^t - q_{KG}^{t-1} = \sum_{i=t-1}^t \gamma \log \frac{\pi_\phi^{KG}(p_i | p_{<i})}{\pi_{ref}(p_i | p_{<i})} \quad (4)$$

where p_i is the triple at step i , and $p_{<i}$ is the KG path of the previous i steps. This reward is used to evaluate the quality of generated KG path steps in the iterative reasoning process (KG Triple Screening and Reconstruction in Section 3.3).

For CoT-PRM, each training data includes three parts: the question Q , true CoT c_w , and false CoT c_l , which is (Q, c_w, c_l) . Among them, the true samples obtain the ground truth, and the false samples are obtained in the following ways: (1) Factual errors (wrong content in reasoning steps); (2) Step skipping (omitting some reasoning steps). (3) Step redundancy (Some steps are unrelated to the QA pairs.). Similar to KG-PRM, we adopt the causal LM log-likelihood ratio as the reward function. For a CoT c and a CoT policy model π_θ^{CoT} , the reward function is:

$$r_{CoT}(c | Q) = \beta \log \frac{\pi_\theta^{CoT}(c | Q)}{\pi_{ref}(c | Q)} \quad (5)$$

where β is used to control the signal strength. For data (Q, c_w, c_l) , the loss function of CoT-PRM is:

$$L_{CoT} = -\log \sigma \left(\beta \log \frac{\pi_\theta^{CoT}(c_w | Q)}{\pi_{ref}(c_w | Q)} - \beta \log \frac{\pi_\theta^{CoT}(c_l | Q)}{\pi_{ref}(c_l | Q)} \right) \quad (6)$$

Then CoT-PRM can distinguish more complete and concise chains and identify factual errors. Similarly, the process rewards at step t for CoT-PRM is,

$$r_{CoT}^t = q_{CoT}^t - q_{CoT}^{t-1} = \sum_{i=t-1}^t \beta \log \frac{\pi_\theta^{CoT}(c_i | c_{<i})}{\pi_{ref}(c_i | c_{<i})} \quad (7)$$

where c_i is the CoT step at step i , and $c_{<i}$ is the previous i CoT steps. In this way, CoT-PRM can be trained with outcome signals and used to evaluate the CoT steps in the iterative reasoning (CoT Step Generation in Section 3.3).

Alternating Optimization After initialization, the two implicit PRMs learn to distinguish between true and false samples. Since the quality of CoT and KG paths is highly correlated, we design a co-training strategy, enabling both implicit PRMs to learn from dual data modalities.

The core of co-training is to enhance the implicit PRMs with data from both modalities. For data (Q, p_w, p_l) from the KG, we convert the true sample p_w and the false sample p_l into continuous texts as the corresponding preference pairs cp_w and cp_l . Then the derived data (Q, cp_w, cp_l) is used to train the CoT-PRM. The loss function L_{CoT-p} still maximizes the probability of the true sample cp_w and minimizes the probability of the false sample cp_l , that is,

$$L_{CoT-p} = -\log \sigma(r_{CoT}(cp_w) - r_{CoT}(cp_l)) \quad (8)$$

For the data from CoT (Q, c_w, c_l), we extract triples in each step and connect them into entity-relation paths to get KG paths pc_w and pc_l . After that, we construct preference signals of KG paths pc_w and pc_l to train the KG-PRM.

$$L_{KG-c} = -\log \sigma(r_{KG}(pc_w) - r_{KG}(pc_l)) \quad (9)$$

Through co-training, the KG-derived CoT preference pairs with structure constraints enable the CoT-PRM to capture topological patterns for long-chain reasoning. Meanwhile, CoT-derived KG preference pairs expand training diversity for KG-PRM by introducing relational paths absent in the original KG. These two PRMs will subsequently serve as evaluators for evaluating KG paths and CoTs.

3.3 Iterative Reasoning Framework

The reasoning process can be divided into four parts: **KG Path and CoT Initialization**, **KG triple screening and reconstruction**, **CoT step generation**, and **Final Answer Generation**. The process is shown in Algorithm 1. We apply KG-PRM and CoT-PRM in the reasoning process to evaluate the KG paths and CoTs. Then they are selected through the Best-of-N (BoN) method based on the process reward.

KG Path and CoT Initialization Before the iteration, we first get the initial KG triple and the first CoT step. Specifically, all-mpnet-base-v2 (Reimers and Gurevych 2019; Song et al. 2020) obtains the embeddings of question Q and all KG triples $tr \in G$. The set tr_1 is composed of the top m similar triples to Q according to the cosine similarity. After that, the LLM retrieves the KG triple p_1 starting from the entities e_1 in question Q from the set tr_1 . The selection is based on the rewards from KG-PRM through the BoN method. Then it will be submitted to the LLM along with q to generate the corresponding CoT step c_1 . c_1 is also screened by CoT-PRM through the BoN method. At this point, the CoT c and the KG path p are initialized with $c = \{c_1\}$ and $p = \{p_1\}$. After that, it goes into the iterative process.

KG Triple Screening and Reconstruction To obtain the KG path at step i , we fetch the next KG triple based on the question Q and the previous CoT step c_{i-1} . First, we concatenate them together to get the string $s_i = Q \oplus c_{i-1}$. After that, s_i will be input into the all-mpnet-base-v2 to obtain the embedding, and select the set tr_i composed of the top m similar triples to it. The LLM retrieves the most relevant triples from tr_i according to the question Q and the previous CoT step c_{i-1} and reconstructs them. Reconstruction refers to the structural adjustment of the selected triples to ensure that the head entities of the triples come from the entities e_{i-1} in the CoT step c_{i-1} , as shown in Figure 3.




Figure 3: Triple reconstruction for entity consistency.

This process obtains the next KG triple p_i based on the rewards from KG-PRM through the BoN method and updates

Algorithm 1: Reasoning Process

Input: question Q , knowledge graph G , max iterations n
Output: Final answer A^*

- 1: $tr_1 \leftarrow \text{top-m}(\text{sim}(Q, tr)), tr \in G$ // tr : KG triple.
- 2: $p_1 \leftarrow \text{LLM_retrieve}(Q, tr_1)$ // p_1 : Initial KG triple.
- 3: $c_1 \leftarrow \text{LLM_generate}(Q, p_1)$ // c_1 : First CoT step.
- 4: $p \leftarrow \{p_1\}, c \leftarrow \{c_1\}, i \leftarrow 2$ // Initialization.
- 5: **while** $i \leq n + 1$ **do**
- 6: $s_i \leftarrow Q \oplus c_{i-1}$ // s_i : concatenation for selecting.
- 7: $tr_i \leftarrow \text{top-m}(\text{sim}(s_i, tr)), tr \in G$
- 8: $p_i \leftarrow \text{LLM_retrieve_reconstruct}(Q, c_{i-1}, tr_i)$ // Get KG triple.
- 9: $p \leftarrow p \cup \{p_i\}$ // Update KG path.
- 10: $c_i \leftarrow \text{LLM_generate}(Q, p_i, c)$ // Get CoT step.
- 11: $c \leftarrow c \cup \{c_i\}$ // Update CoT.
- 12: **if** "answer" **in** c_i **then**
- 13: **break** // "answer" is a keyword.
- 14: **end if**
- 15: $i \leftarrow i + 1$
- 16: **end while**
- 17: $\text{prompt}_{hard} \leftarrow [Q, c, p]$ // Plain text prompt.
- 18: $\text{prompt}_{soft} \leftarrow \text{text}(p)$ // Graph structure prompt.
- 19: $A^* \leftarrow \text{LLM_generate}(\text{prompt}_{hard}, \text{prompt}_{soft})$
- 20: **return** A^*

the current KG path to $p = \{p_1, \dots, p_i\}$. This triple p_i will guide the expansion of CoT in the next part.

CoT Step Generation We generate the next CoT step based on the KG triple p_i , and the entities in it will guide the selection of the next KG triple. We concatenate the question Q , p_i and the previous CoT $c = \{c_1, \dots, c_{i-1}\}$ according to the prompt template. Then we use the prompt to generate the next CoT step. This process obtains the next CoT step c_i based on the rewards from CoT-PRM through the BoN method and updates the current CoT to $c = \{c_1, \dots, c_i\}$.

The iteration stops when CoT c contains the keyword "answer" or reaches the preset limit. The CoT c and the KG path p will support the final answer generation.

Final Answer Generation In generation, the prompts submitted to the LLM are divided into hard prompts (plain text) and soft prompts (textualized graph) (Hu et al. 2025; He et al. 2025). They are submitted to the LLM to generate the final answer A^* . By combining them, text and graph information can both be preserved to reduce knowledge loss.

4 Experiments

4.1 Experiment Setup

Datasets. Following previous works (Luo et al. 2024a; Wang et al. 2023a; Luo et al. 2024c), we conduct experiments on two datasets, WebQuestionSP (WebQSP) (Yih et al. 2016) and Complex WebQuestions (CWQ) (Talmor and Berant 2018).

Baselines. We compare DPRM with 13 baselines across 4 categories: (1) LLM only, including Qwen2-7B (Yang et al. 2024), Llama-2-7B (Touvron et al. 2023), and Llama-3.1-8B

Category	Method	WebQSP		CWQ	
		F1 Score	Hit@1	F1 Score	Hit@1
LLM only	Qwen2-7B (Yang et al. 2024)	0.3550	0.5080	0.2160	0.2530
	Llama-2-7B (Touvron et al. 2023)	0.3650	0.5640	0.2140	0.2840
	Llama-3.1-8B (Dubey et al. 2024)	0.3480	0.5550	0.2240	0.2810
KG+LLM	G-Retriever (He et al. 2025)	0.4674	0.6808	0.3396	0.4721
	GRAG (Hu et al. 2025)	0.5022	0.7236	0.3649	0.5018
	SubgraphRAG (Li, Miao, and Li 2025)	0.7057	<u>0.8661</u>	0.4716	0.5698
	RoG (Luo et al. 2024a)	0.7080	<u>0.8570</u>	0.5620	<u>0.6260</u>
	GNN-RAG (Mavromatis and Karypis 2024)	<u>0.7130</u>	0.8060	<u>0.5940</u>	<u>0.6170</u>
ORM+LLM	SkyworkRM-Llama3.1-8B (Liu et al. 2024)	0.5751	0.7329	0.3465	0.5755
	ArmoRM-Llama3-8B (Wang et al. 2024)	0.5816	0.7198	0.3345	0.5531
PRM+LLM	RLHFlow-8B-Mistral-Data (Xiong et al. 2024)	0.5863	0.7320	0.3431	0.5540
	RLHFlow-8B-DeepSeek-Data (Dong et al. 2024a)	0.5356	0.6815	0.3572	0.5718
	ImplicitPRM-DPO (Yuan et al. 2024)	0.5744	0.7266	0.3538	0.5650
Our method	DPRM+Llama-3.1-8B	0.7228	0.8715	0.5988	0.7301
	DPRM+Llama-2-7B	0.7220	0.8718	0.5886	0.7299

Table 1: Model performance on two datasets comparing five categories of methods. The best results are **bolded**, and the best results in baselines are underlined.

(Dubey et al. 2024). (2) KG+LLM, including G-Retriever (He et al. 2025), GRAG (Hu et al. 2025), SubgraphRAG (Li, Miao, and Li 2025), RoG (Luo et al. 2024a), and GNN-RAG (Mavromatis and Karypis 2024). (3) ORM+LLM, including SkyworkRM-Llama3.1-8B (Liu et al. 2024) and ArmoRM-Llama3-8B (Wang et al. 2024). (4) PRM+LLM, including RLHFlow-8B-Mistral-Data (Xiong et al. 2024), RLHFlow-8B-DeepSeek-Data (Dong et al. 2024a), and ImplicitPRM-DPO (Yuan et al. 2024).

Evaluation Metrics. Following previous works (Luo et al. 2024a; Hu et al. 2025; Luo et al. 2024c), we use Hit@1 and the F1 score as evaluation metrics. Hit@1 checks if the ground truth exists in the generated answers. The F1 score is a harmonic average of accuracy and recall, providing a metric that balances false positives and false negatives.

Implementations. We train PRMs based on Llama-3.1-8B-Instruct with $\beta = 0.05$ and $\gamma = 0.05$, which are empirically determined. We select the number of samples $N = 8$ for screening and the number of top-m related triples $m = 25$. The maximum number n of iterations is set to 4, considering that the data is at most 4 hops.

4.2 Main Results

We compare our method, DPRM, to other baselines on the datasets. As Table 1 shows, DPRM performs best on both datasets. The F1 score on WebQSP and CWQ is 1.4% and 0.8% above the best baseline, and Hit@1 is 0.7% and 16.6% above. It shows that DPRM can effectively enhance the reasoning ability of the LLM.

The overall performance of KG+LLM baselines is better than that of LLM only baselines, indicating that KGs are important in MHQA tasks. It is also found that RMs can effectively improve the reasoning ability of LLMs in MHQA

tasks. The ORM+LLM baselines increase the F1 score of the corresponding LLM by up to 67.1% and Hit@1 by up to 104.8%. Similarly, the PRM+LLM baselines increase the F1 score by up to 68.5% and Hit@1 by up to 103.5%. However, since these traditional RMs fail to consider the structural constraints of KGs, their overall performance is worse than that of KG+LLM baselines. This again highlights the importance of KG. In addition, because MHQA tasks involve fewer reasoning steps than mathematical reasoning tasks, the performance gap between ORMs and PRMs is not very big. On the webqsp dataset where questions are 1-2 hops, the worst-performing RM is even a PRM (RLHFlow-8B-DS-Data). This indicates that fewer reasoning steps in MHQA tasks can limit the advantage of PRMs.

We also observe that larger LLMs do not always perform better than smaller LLMs. For the performance of Llama-2-7B and Llama-3.1-8B, Llama-2-7B has a higher F1 and Hit@1 on the WebQSP dataset, by 4.9% and 1.6%. In the CWQ dataset, its Hit@1 is still higher. This suggests that increasing the parameters does not inherently enhance the graph reasoning ability of LLMs.

4.3 Ablation Study

Method	WebQSP		CWQ	
	F1 Score	Hit@1	F1 Score	Hit@1
DPRM	0.7228	0.8715	0.5988	0.7301
w/o Co-training	0.7009	0.8581	0.5419	0.6824
w/o Iteration	0.7095	0.8587	0.5824	0.7234
w/o both	0.6842	0.8500	0.5266	0.6692

Table 2: Performances of three model variables.

We conduct a series of evaluations of DPRM to see which component plays a key role, including removing co-training in training, removing iteration in reasoning, and removing both of the two parts. We can see the performances of variables all decrease, as shown in Table 2. It indicates that every component is indispensable. Among them, removing co-training drops model performance more than removing iteration. On the CWQ dataset, F1 of the former decreased by 6.8%, and Hit@1 decreased by 5.5% more than that of the latter. The situation is similar on WebQSP. This suggests that co-training plays a more central role in reasoning. It enables the CoT-PRM to capture topological patterns for long-chain reasoning and expand training diversity for KG-PRM by introducing relational paths absent in the original KG.

4.4 Analytical Experiments




Figure 4: Performances on the CoT-PRMs trained on original and KG-derived data.

Topological Pattern Learning for CoT-PRM To test whether the data derived from KG paths allows CoT-PRM to learn topological patterns, we randomly generate 800 new KG preference pairs and use the CoT-PRMs trained on original and KG-derived data to evaluate them. It tests whether the two CoT-PRMs can distinguish true and false samples of structured data. As shown in Figure 4, the accuracy of the CoT-PRM trained on KG-derived data is higher than that of the original-data-trained CoT-PRM on both datasets. It indicates that co-training enables the CoT-PRM to capture topological patterns for long-chain reasoning.




Figure 5: Proportion of Same and Different Entities/Relations in CoT-Derived KG Data.

New Relational Paths for KG-PRM To evaluate the impact of co-training on the training data for KG-PRM, we calculate the proportion of entities and relations in the CoT-derived KG that are the same as and different from those in the original KG. As shown in Figure 5, it keeps some key entities but shows big differences in relations. Specifically, the CoT-derived KG data

does not follow standard KG relation formats like “celebrities.romantic_relationship.celebrity”, but instead uses more accurate terms like “marry”. It indicates that CoT-derived KG data expand training diversity for KG-PRM by introducing paths absent in the original KG.

Method	WebQSP		CWQ	
	F1 Score	Hit@1	F1 Score	Hit@1
SkyworkRM	0.5751	0.7329	0.3465	0.5755
ArmoRM	0.5816	0.7198	0.3345	0.5531
RLHFlow-Mistral	0.5863	0.7320	0.3431	0.5540
RLHFlow-DeepSeek	0.5356	0.6815	0.3572	0.5718
ImplicitPRM-DPO	0.5744	0.7266	0.3538	0.5650
KG-PRM (Our method)	0.5946	0.7346	0.4283	0.6341

Table 3: Performances of RMs.

KG Structural Significance As shown in Table 3, our KG-PRM outperforms other RMs, especially on CWQ, which has more multi-hop questions. This shows that our method can make KG-PRM consider the graph structure constraint in KGs and learn to identify accurate KG paths that follow KG logic rules. In addition, in MHQA tasks, structural knowledge is essential, which can significantly enhance the LLM reasoning.

Method	WebQSP		CWQ	
	F1 Score	Hit@1	F1 Score	Hit@1
DPRM (full data)	0.7228	0.8715	0.5988	0.7301
DPRM (partial data)	0.6670	0.8089	0.5843	0.7190

Table 4: Performances with Full/Partial Training Data.

Partial Data Robustness We evaluate DPRM with CoT-PRM and KG-PRM trained on partial datasets. Each PRM is fine-tuned on 1,000 preference pairs of CoTs or KG paths, respectively. As shown in Table 4, DPRM achieves 92.8% of the full-data performance using only one-eighth of the training data. Compared with the results in Table 1, our DPRM trained on partial data still outperforms 10 baselines.

5 Conclusion

In this paper, we propose a dual implicit process reward model (DPRM) for automatic annotation and collaborative reasoning in MHQA tasks. It trains KG-PRM and CoT-PRM to assign rewards for CoT and KG path steps during reasoning without using step-wise annotations. Different from existing implicit PRMs, KG-PRM is designed for capturing the structural information of KG; co-training of both PRMs enhances the consistency between CoT and KG: CoT-PRM learns to capture the topological patterns in KG, and KG-PRM learns from more relational and diverse paths. The experiment results support the above statements and show that DPRM achieves state-of-the-art (SOTA) performance among 13 baselines.

Acknowledgments

This paper is partially supported by National Natural Science Foundation of China (NSFC Grant No. 62576353 and No. 62306330) and National Natural Science Foundation of China Distinguished Young Scholar Project (Grant No. 62125604).

References

- Cao, M.; Shu, L.; Yu, L.; Zhu, Y.; Wichers, N.; Liu, Y.; and Meng, L. 2024. Beyond Sparse Rewards: Enhancing Reinforcement Learning with Language Model Critique in Text Generation. *arXiv:2401.07382*.
- Chan, A. J.; Sun, H.; Holt, S.; and van der Schaar, M. 2024. Dense Reward for Free in Reinforcement Learning from Human Feedback. *arXiv:2402.00782*.
- Christiano, P. F.; Leike, J.; Brown, T.; Martic, M.; Legg, S.; and Amodei, D. 2017. Deep reinforcement learning from human preferences. *Advances in neural information processing systems*, 30.
- Cui, G.; Yuan, L.; Wang, Z.; Wang, H.; Li, W.; He, B.; Fan, Y.; Yu, T.; Xu, Q.; Chen, W.; Yuan, J.; Chen, H.; Zhang, K.; Lv, X.; Wang, S.; Yao, Y.; Han, X.; Peng, H.; Cheng, Y.; Liu, Z.; Sun, M.; Zhou, B.; and Ding, N. 2025. Process Reinforcement through Implicit Rewards. *arXiv:2502.01456*.
- Dong, H.; Xiong, W.; Pang, B.; Wang, H.; Zhao, H.; Zhou, Y.; Jiang, N.; Sahoo, D.; Xiong, C.; and Zhang, T. 2024a. RLHF Workflow: From Reward Modeling to Online RLHF. *CoRR*, abs/2405.07863.
- Dong, Z.; Peng, B.; Wang, Y.; Fu, J.; Wang, X.; Shan, Y.; and Zhou, X. 2024b. EffiQA: Efficient Question-Answering with Strategic Multi-Model Collaboration on Knowledge Graphs. *arXiv preprint arXiv:2406.01238*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; Metropolitan, D.; Ness, R. O.; and Larson, J. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Guan, X.; Zhang, L. L.; Liu, Y.; Shang, N.; Sun, Y.; Zhu, Y.; Yang, F.; and Yang, M. 2025. rStar-Math: Small LLMs Can Master Math Reasoning with Self-Evolved Deep Thinking. *arXiv preprint arXiv:2501.04519*.
- He, X.; Tian, Y.; Sun, Y.; Chawla, N.; Laurent, T.; LeCun, Y.; Bresson, X.; and Hooi, B. 2025. G-retriever: Retrieval-augmented generation for textual graph understanding and question answering. *Advances in Neural Information Processing Systems*, 37: 132876–132907.
- Hu, Y.; Lei, Z.; Zhang, Z.; Pan, B.; Ling, C.; and Zhao, L. 2025. GRAG: Graph Retrieval-Augmented Generation. *arXiv:2405.16506*.
- Jin, B.; Xie, C.; Zhang, J.; Roy, K. K.; Zhang, Y.; Li, Z.; Li, R.; Tang, X.; Wang, S.; Meng, Y.; et al. 2024. Graph chain-of-thought: Augmenting large language models by reasoning on graphs. *arXiv preprint arXiv:2404.07103*.
- Li, M.; Miao, S.; and Li, P. 2025. Simple Is Effective: The Roles of Graphs and Large Language Models in Knowledge-Graph-Based Retrieval-Augmented Generation. *arXiv:2410.20724*.
- Li, S.; He, Y.; Guo, H.; Bu, X.; Bai, G.; Liu, J.; Liu, J.; Qu, X.; Li, Y.; Ouyang, W.; et al. 2024. GraphReader: Building Graph-based Agent to Enhance Long-Context Abilities of Large Language Models. *arXiv preprint arXiv:2406.14550*.
- Li, W.; Yang, W.; Hou, Y.; Liu, L.; Liu, Y.; and Li, X. 2025. SARATR-X: Toward Building a Foundation Model for SAR Target Recognition. *IEEE Transactions on Image Processing*, 34(1): 869–884.
- Lightman, H.; Kosaraju, V.; Burda, Y.; Edwards, H.; Baker, B.; Lee, T.; Leike, J.; Schulman, J.; Sutskever, I.; and Cobbe, K. 2023. Let’s verify step by step. In *The Twelfth International Conference on Learning Representations*.
- Liu, C. Y.; Zeng, L.; Liu, J.; Yan, R.; He, J.; Wang, C.; Yan, S.; Liu, Y.; and Zhou, Y. 2024. Skywork-Reward: Bag of Tricks for Reward Modeling in LLMs. *arXiv:2410.18451*.
- Liu, L.; Sun, S.; Zhi, S.; Shi, F.; Liu, Z.; Heikkilä, J.; and Liu, Y. 2025. A Causal Adjustment Module for Debiasing Scene Graph Generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(5): 4024–4043.
- Lu, J.; Dou, Z.; Wang, H.; Cao, Z.; Dai, J.; Feng, Y.; and Guo, Z. 2024. Autopsv: Automated process-supervised verifier. *Advances in Neural Information Processing Systems*, 37: 79935–79962.
- Luo, L.; Li, Y.-F.; Haffari, G.; and Pan, S. 2024a. Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning. In *International Conference on Learning Representations*.
- Luo, L.; Liu, Y.; Liu, R.; Phatale, S.; Guo, M.; Lara, H.; Li, Y.; Shu, L.; Zhu, Y.; Meng, L.; Sun, J.; and Rastogi, A. 2024b. Improve Mathematical Reasoning in Language Models by Automated Process Supervision. *arXiv:2406.06592*.
- Luo, L.; Zhao, Z.; Gong, C.; Haffari, G.; and Pan, S. 2024c. Graph-constrained reasoning: Faithful reasoning on knowledge graphs with large language models. *arXiv preprint arXiv:2410.13080*.
- Mavromatis, C.; and Karypis, G. 2024. GNN-RAG: Graph Neural Retrieval for Large Language Model Reasoning. *arXiv preprint arXiv:2405.20139*.
- Nishat, N. A. Z.; Coletta, A.; Bellomarini, L.; Amouzouvi, K.; Lehmann, J.; and Vahdati, S. 2025. Aligning Knowledge Graphs and Language Models for Factual Accuracy. *arXiv:2507.13411*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in neural information processing systems*, 36: 53728–53741.

- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3982–3992.
- Sanmartin, D. 2024. KG-RAG: Bridging the Gap Between Knowledge and Creativity. *arXiv preprint arXiv:2405.12035*.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2020. MpNet: Masked and permuted pre-training for language understanding. *Advances in neural information processing systems*, 33: 16857–16867.
- Talmor, A.; and Berant, J. 2018. The Web as a Knowledge-Base for Answering Complex Questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 641–651.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, H.; Xiong, W.; Xie, T.; Zhao, H.; and Zhang, T. 2024. Interpretable Preferences via Multi-Objective Reward Modeling and Mixture-of-Experts. In *EMNLP (Findings)*, 10582–10592.
- Wang, K.; Duan, F.; Wang, S.; Li, P.; Xian, Y.; Yin, C.; Rong, W.; and Xiong, Z. 2023a. Knowledge-driven cot: Exploring faithful reasoning in llms for knowledge-intensive question answering. *arXiv preprint arXiv:2308.13259*.
- Wang, P.; Li, L.; Shao, Z.; Xu, R.; Dai, D.; Li, Y.; Chen, D.; Wu, Y.; and Sui, Z. 2023b. Math-shepherd: Verify and reinforce llms step-by-step without human annotations. *arXiv preprint arXiv:2312.08935*.
- Wang, X.; Song, Y.; Liu, C.; Luo, T.; Liu, B.; Xie, Z.; and Huang, M. 2025. DCMKC: A Dual Consistency Matching Approach for Multi-hop Question Answering in LLMs. In Christodoulopoulos, C.; Chakraborty, T.; Rose, C.; and Peng, V., eds., *Findings of the Association for Computational Linguistics: EMNLP 2025*, 259–273. Suzhou, China: Association for Computational Linguistics. ISBN 979-8-89176-335-7.
- Wu, J.; Deng, W.; Li, X.; Liu, S.; Mi, T.; Peng, Y.; Xu, Z.; Liu, Y.; Cho, H.; Choi, C.-I.; Cao, Y.; Ren, H.; Li, X.; Li, X.; and Zhou, Y. 2025. MedReason: Eliciting Factual Medical Reasoning Steps in LLMs via Knowledge Graphs. *arXiv:2504.00993*.
- Xiong, W.; Dong, H.; Ye, C.; Wang, Z.; Zhong, H.; Ji, H.; Jiang, N.; and Zhang, T. 2024. Iterative Preference Learning from Human Feedback: Bridging Theory and Practice for RLHF under KL-Constraint. *arXiv:2312.11456*.
- Xu, Z.; Cruz, M. J.; Guevara, M.; Wang, T.; Deshpande, M.; Wang, X.; and Li, Z. 2024. Retrieval-Augmented Generation with Knowledge Graphs for Customer Service Question Answering. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2024*, 2905–2909. ACM.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; Wang, J.; Yang, J.; Tu, J.; Zhang, J.; Ma, J.; Yang, J.; Xu, J.; Zhou, J.; Bai, J.; He, J.; Lin, J.; Dang, K.; Lu, K.; Chen, K.; Yang, K.; Li, M.; Xue, M.; Ni, N.; Zhang, P.; Wang, P.; Peng, R.; Men, R.; Gao, R.; Lin, R.; Wang, S.; Bai, S.; Tan, S.; Zhu, T.; Li, T.; Liu, T.; Ge, W.; Deng, X.; Zhou, X.; Ren, X.; Zhang, X.; Wei, X.; Ren, X.; Liu, X.; Fan, Y.; Yao, Y.; Zhang, Y.; Wan, Y.; Chu, Y.; Liu, Y.; Cui, Z.; Zhang, Z.; Guo, Z.; and Fan, Z. 2024. Qwen2 Technical Report. *arXiv:2407.10671*.
- Yih, W.-t.; Richardson, M.; Meek, C.; Chang, M.-W.; and Suh, J. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 201–206.
- Yuan, L.; Li, W.; Chen, H.; Cui, G.; Ding, N.; Zhang, K.; Zhou, B.; Liu, Z.; and Peng, H. 2024. Free Process Rewards without Process Labels. *arXiv:2412.01981*.
- Zhang, D.; Zhoubian, S.; Hu, Z.; Yue, Y.; Dong, Y.; and Tang, J. 2024. Rest-mcts*: Llm self-training via process reward guided tree search. *Advances in Neural Information Processing Systems*, 37: 64735–64772.
- Zhao, J.; Liu, R.; Zhang, K.; Zhou, Z.; Gao, J.; Li, D.; Lyu, J.; Qian, Z.; Qi, B.; Li, X.; and Zhou, B. 2025. GenPRM: Scaling Test-Time Compute of Process Reward Models via Generative Reasoning. *arXiv:2504.00891*.
- Zhao, X.; Blum, M.; Yang, R.; Yang, B.; Carpintero, L. M.; Pina-Navarro, M.; Wang, T.; Li, X.; Li, H.; Fu, Y.; Wang, R.; Zhang, J.; and Li, I. 2024. AGENTiGraph: An Interactive Knowledge Graph Platform for LLM-based Chatbots Utilizing Private Data. *arXiv:2410.11531*.
- Ziyi, H.; Jun, L.; Zhongzhi, L.; Yuzhong, L.; Zheng, X.; and Yiping, S. 2025. RMath: A Logic Reasoning-Focused Datasets Toward Mathematical Multistep Reasoning Tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 24104–24112.