

PretrainRL: Alleviating Factuality Hallucination of Large Language Models at the Beginning

Langming Liu^{1*}, Kangtao Lv^{1,2*}, Haibin Chen¹, Weidong Zhang¹, Yejing Wang³, Shilei Liu¹, Xin Tong¹, Yujin Yuan¹, Yongwei Wang^{2,4}, Wenbo Su¹, Bo Zheng^{1†}

¹Future Living Lab of Alibaba ²Zhejiang University

³Alibaba Group ⁴Shanghai AI Laboratory

{liulangming.llm, lvkangtao.lkt, bozheng}@alibaba-inc.com

Abstract

Large language models (LLMs), despite their powerful capabilities, suffer from factual hallucinations where they generate verifiable falsehoods. We identify a root of this issue: the imbalanced data distribution in the pretraining corpus, which leads to a state of "low-probability truth" and "high-probability falsehood". Recent approaches, such as teaching models to say "I don't know" or post-hoc knowledge editing, either evade the problem or face catastrophic forgetting. To address this issue from its root, we propose **PretrainRL**, a novel framework that integrates reinforcement learning into the pretraining phase to consolidate factual knowledge. The core principle of PretrainRL is "**debiasing then learning**." It actively reshapes the model's probability distribution by down-weighting high-probability falsehoods, thereby making "room" for low-probability truths to be learned effectively. To enable this, we design an efficient negative sampling strategy to discover these high-probability falsehoods and introduce novel metrics to evaluate the model's probabilistic state concerning factual knowledge. Extensive experiments on three public benchmarks demonstrate that PretrainRL significantly alleviates factual hallucinations and outperforms state-of-the-art methods.

1 Introduction

LLMs (OpenAI, 2024; DeepSeek-AI, 2024; Grattafiori and et al., 2024; Yang et al., 2025) have recently drawn significant attention for their strong capabilities in natural language processing (NLP). Although LLMs show powerful knowledge retention, LLMs suffer from ubiquitous and difficult-to-cure hallucination issues. The taxonomy of such hallucinations is constantly evolving, generally categorized as factuality-related, instruction-related, and logic-related errors (Ji et al., 2023; Tonmoy

*Equal contribution

†Corresponding author




Figure 1: Overall comparison on POPQA

et al., 2024; Zhang et al., 2025). In this work, we focus on **factual hallucinations**, where model generations contain clear and verifiable errors. This type of hallucination is arguably the most fundamental and has the most extensive impact, since it pertains to statements that can be easily disproven with ground-truth knowledge. For example, when the user ask "What is the capital of Arno?", and LLM response is that "Rome", however the answer is "Florence".

Several efforts have been made to address or alleviate factual hallucinations. A prevalent approach is to teach LLMs to say "**I don't know**" (IDK) when uncertain (Lin et al., 2022; Cao, 2023; Manakul et al., 2023). This involves learning a preference for refusing to answer when generation probabilities are low, implemented through preference alignment, novel decoding methods, or specialized loss functions. While this direction is safe and effective at avoiding hazardous responses, it has notable drawbacks: models can become "lazy" by over-producing IDK responses, which in turn restricts generative diversity (Kadavath et al., 2022; Touvron et al., 2023; Bai et al., 2022). Crucially, this method cannot directly resolve factual hallucinations, but merely conceals them. Another line of research, known as **knowledge editing** (Cao et al., 2021; Meng et al., 2022a,b; Zheng et al., 2023; Wang et al., 2024c), aims to efficiently and pre-

cisely correct factual errors. This involves locating and making subtle modifications to the parameters associated with a specific fact. However, the core challenge lies in its scalability; a large number of edits can degrade overall capabilities and lead to catastrophic forgetting (Mitchell et al., 2021).

Although teaching models to say IDK does not address factual hallucinations at its root, it points to a core source: truths have a low generation probability. More accurately, **low-probability truth** and **high-probability falsehood** jointly cause hallucinations. This phenomenon exists widely in LLMs, and the in-depth reason lies not in noisy or polluted data, but in the pretraining paradigm itself (Xu et al., 2024; Kalai et al., 2025). The next-token prediction (NTP) objective forces the model to fit the corpus data distribution, which is essentially a density estimation (Radford et al., 2018). However, an imbalanced data distribution is inevitable (Gao et al., 2020; Roberts et al., 2020). For certain knowledge domains, this imbalance is extreme: a small amount of “head knowledge” occupies a high cumulative density. In this scenario, head knowledge can interfere with the memorization of “tail knowledge” during training, causing “head falsehoods” (i.e., high-probability falsehoods). For instance, if a corpus contains factual knowledge of “Brand A shoes are red” 1,000 times more frequently than “Brand B shoes are blue,” an LLM may learn a shortcut to associate “shoes are” with “red,” degrading from learning a conditional probability to a marginal one and leading to severe hallucinations. This kind of degradation can also occur for more common data distributions.

The critical challenge is to consolidate this tail knowledge without disrupting existing knowledge. Modifying the corpus distribution is ill-advised as it can introduce new biases. Post-training is also inappropriate, as the model’s knowledge is more “rigid” at this stage, and aggressive knowledge injection can cause catastrophic forgetting or pattern collapse. Therefore, we propose **Pre-training Reinforcement Learning**, a framework that consolidates knowledge during the pretraining stage. The core principle of PretrainRL is “debiasing then learning”: first lowering the probability of falsehoods to “make room,” then increasing the probability of truths. Unlike simply continual pretraining (Wang et al., 2024a; Shi et al., 2024; Ke et al., 2023), which focuses only on learning, this debiasing step is crucial given the limited capacity of parametric knowledge storage. We leverage di-

rect preference optimization (DPO) (Rafailov et al., 2023) for this “probability reshaping.” The main challenge then becomes discovering the rejected samples (i.e., head falsehoods) for DPO, which determines the room for truth learning. We design an efficient sampling approach to obtain high-quality rejected samples for DPO. Furthermore, we introduce three metrics to better monitor and evaluate the state of the model’s generation probability. The contributions of our work are three-fold as follows:

- We reveal that the core source of factual hallucination is the imbalanced distribution in the pre-training corpus and propose PretrainRL, which novelly introduces an RL method to alleviate factual hallucinations during the pretraining phase.
- We propose an efficient and effective method for negative sampling, boosting DPO’s effectiveness on probability reshaping. We also introduce crucial metrics for a comprehensive evaluation of the model’s probabilistic state.
- Thorough and extensive experiments on three public benchmarks demonstrate the superior performance of the proposed PretrainRL compared to other state-of-the-art LLMs. Further analysis provides insights into the study of hallucination.

2 Preliminary

2.1 Knowledge Graph

A knowledge graph (KG) is a structured representation of factual knowledge, where entities (nodes) are interconnected by relations (edges) (Hogan et al., 2021; Chen et al., 2020). KGs serve as a foundational resource for a wide range of NLP tasks by providing explicit, verifiable world knowledge. The basic unit of a KG is the **knowledge triple**, a structured assertion of a fact in the form of (s, p, o) . Here, s is the **subject** (i.e., head entity), o is the **object** (i.e., tail entity), and r is the **predicate** (i.e., relation that connects them). For example, the factual statement “Paris is the capital of France” can be represented as the triple (Paris, capitalOf, France). The subject Paris and object France are entities, while capitalOf is the relation describing their semantic connection. This triple-based structure allows for the systematic organization of complex information. In the context of LLMs, KGs offer a crucial, external source of ground-truth information, providing a pathway to verify model-generated statements and mitigate factual hallucinations.

2.2 Next Token Prediction

The predominant pretraining objective for autoregressive LLMs is **Next-Token Prediction (NTP)**. Given a sequence of tokens $x = \{x_1, x_2, \dots, x_{i-1}\}$, the model is trained to predict the next token x_i by maximizing its conditional probability. Formally, the objective is to minimize the negative log-likelihood (NLL) over a large-scale pretraining corpus \mathcal{D} :

$$\mathcal{L}_{\text{NTP}} = -\mathbb{E}_{x \sim \mathcal{D}} \sum_{i=1}^{|x|} \log P(x_i | x_{<i}; \theta) \quad (1)$$

where x is a sequence in the corpus, $x_{<i}$ denotes the preceding tokens, and θ represents the model parameters. This self-supervised objective compels the base model to internalize a large amount of syntactic, semantic, and factual knowledge from the training data in order to make accurate predictions. By fitting the data distribution of the pretraining corpus (i.e., density estimation), the model essentially learns a statistical representation of the world. However, the model’s learned distribution is susceptible to biases present in the data, such as the frequency imbalance of factualities, which is a primary source of hallucination.

3 Methodology

We first analyze the source of the factual hallucination of LLMs, then provide the solution to address this issue.

3.1 Source of Factual Hallucination

The KG triple can be considered to be the basic unit of the factuality in corpus. Suppose that the triple (s, p, o) is a factual knowledge in sequence x . Then, minimizing NLL loss can also maximize

$$P(o|s, p; \theta), \quad (2)$$

which is a conditional probability of the object. As in the previous statement, the critical source of factual hallucinations is low-probability truth and high-probability falsehood. To put it more deeply, the imbalanced distribution of knowledge causes such probability error of LLMs.

Given a certain knowledge category, suppose that this imbalanced data structure consists of only two types of triples, $D_1 = (s_1, p, o_1)$, $D_2 = (s_2, p, o_2)$. The frequency of D_1 far exceeds that of D_2 , which means that frequency ξ_1 may be hundreds of frequency ξ_2 in the corpus. We call D_1 , D_2

head knowledge and **tail knowledge**, respectively. The NLL loss is then partially maximizing the weighted sum of conditional probability of head and tail knowledge, formed as

$$\xi_1 P(o_1 | s_1, p; \theta) + \xi_2 P(o_2 | s_2, p; \theta), \quad (3)$$

where ξ_1, ξ_2 is the frequency, and $\xi_1 \gg \xi_2$. As such an imbalanced knowledge exists widely in corpus, the pretraining task (i.e., NTP) may force LLMs to take shortcuts to obtain this objective. Specifically, LLMs tend to neglect the subject s of the context, instead degrading to learn a margin $P(o|p; \theta)$ of conditional probability. In this case, LLM have a high probability $P(o_1 | p; \theta)$ to predict object o_1 , regardless of given context s_1 or s_2 . Consequently, head knowledge D_1 successfully permeates tail knowledge D_2 , leading to severe factual hallucination. From the perspective of the attention mechanism, LLMs learn high attention scores on predicate p , and low attention scores on subject s . This is the reason that, even though we emphasize the subject in the prompt, LLMs will persist in outputting a popular answer rather than the correct answer.

3.2 Negative Sampling

The previous theoretical analysis points out that head knowledge D_1 squeezes the memory room of tail knowledge D_2 , resulting in very low conditional probability $P(o_2 | s, p; \theta)$ so as to degenerate to a marginal probability. To address this issue, the primary task is to locate those head knowledge (i.e., negative samples) that affect LLMs’ memory of the tail and need to be debiased. The most intuitive way is to classify knowledge into multiple categories, and then use the popularity or frequency of the knowledge in each category to obtain the quantiles of knowledge. The quantiles decide whether the knowledge belongs to head. However, the pretraining corpus is usually not available for us. Even if we can access the corpus, calculating popularity of knowledge is complicated, since the forms of text are irregular.

Therefore, we propose an efficient sampling method to obtain negative samples, using the **beam search** decoding method. Specifically, given a knowledge triple (s, p, o) , we use a prompt consisting of s, p and ask LLMs to continue (see Appendix A for prompt details), where we use beam search to let LLM decode several answers. Then we obtain the probability of each answer using the

probability of cumulative multiplication of all tokens. Notice that ground truth o should be filtered out. We order the answers with probability and get the negative samples using a threshold or a preset quantity (e.g., 5, 10). The negative samples obtained directly reflect the LLM’s inherent bias for head knowledge, which are actually our targets.

Furthermore, since models inherently prioritize generating head knowledge for a given question category, a subset of the full dataset suffices to statistically approximate the head knowledge distribution. Taking EntityQuestions (170k instances) as an example, computing statistics over the entire dataset is computationally redundant. We sample 1k instances per **question category** and count the top 20 high-frequency object, then compare this distribution with the top 20 obtained from the full dataset. The cosine similarity between the two frequency distributions is 0.88, and the Spearman rank correlation is 0.85, indicating that the sampled distribution closely matches the full-data distribution and confirming the effectiveness of our approach.

3.3 Optimization: PretrainRL

Our target is to reshape the model’s probabilities of knowledge to avoid factual hallucinations. As is well known, RL is designated for lossless probability reshaping. In addition, selecting the specific optimization approach and timing is also crucial.

3.3.1 Methods: DPO

We can easily obtain the negative samples (i.e., using sampling methods as mentioned in Section 3.2) and positive samples (i.e., the ground truth) of knowledge. Therefore, the natural approach is to use direct preference optimization (DPO) for probability reshaping. As factual knowledge is verifiable, reinforcement learning with verifiable rewards (RLVR) methods (Lambert et al., 2024; Shao et al., 2024) is an alternative. However, RLVR is inefficient in the pretraining stage, as the amount of knowledge is gigantic. Moreover, DPO matches more closely with our proposed efficient negative sampling methods in Section 3.2. Thus, we choose DPO as the optimization method to mitigate factual hallucinations.

3.3.2 Timing: Pretraining

RL or SFT methods are typically applied in the post-training stage. However, as discussed in the introduction, post-training often serves preference learning, rather than large-scale knowledge injection.

Since the posttrained model is resistant to knowledge injection (Table 2), decreasing the upper bound. Moreover, intensive knowledge learning in the posttraining phase can lead to catastrophic forgetting, as demonstrated in Section 4.5. In contrast, pretrained models contact massive and diverse data, and possess unparalleled breadth and coverage of knowledge. Moreover, pretrained models have better plasticity, making it easier for us to perform large-scale probability reshaping (Dong et al., 2025). So we select the pretraining stage to perform RL methods for debiasing. Consequently, this approach is named **PretrainRL**. Specifically, we implement PretrainRL during the continual training (CT) phase.

3.3.3 Objective

For a KG triple (s, p, o) , our goal is to unlearn negative samples (i.e., debiasing the falsehood) and train positive samples (i.e., learning the truth). We apply general DPO loss by masking the prompt $x_{s,p}$ (i.e., synthesized from s, p), focusing on learning the object o . The DPO loss can be formed as follows:

$$\mathcal{L}_{\text{DPO}} = -\mathbb{E}_{(s,p,o) \sim \mathcal{D}} \left(\beta \log \frac{\pi_{\theta}(o_w | x_{s,p})}{\pi_{\text{ref}}(o_w | x_{s,p})} - \beta \log \frac{\pi_{\theta}(o_l | x_{s,p})}{\pi_{\text{ref}}(o_l | x_{s,p})} \right), \quad (4)$$

where π_{θ} is the policy model being optimized, while π_{ref} is a frozen reference model (i.e., a pretrained model). The hyperparameter β controls the strength. This objective encourages the model π_{θ} to increase the relative log-probability of the winner response o_w (i.e., ground truth o) compared to the loser response o_l (i.e., negative sample related to o).

However, two challenges restrict the effectiveness of DPO in probability reshaping. First, when the conditional probability $P(o_w | s, p; \theta)$ is very low, DPO forces $P(o_l | s, p; \theta)$ to be lower, instead of increasing $P(o_w | s, p; \theta)$. Second, simply applying DPO is not coordinated with the pretraining paradigm, raising the risk of pattern collapse. To maintain the model’s core language capabilities, meanwhile, enhancing DPO’s effects to boost probability reshaping, we incorporate a standard NLL loss (Dubey et al., 2024) (i.e., CT loss). It is formulated as:

$$\mathcal{L}_{\text{CT}} = -\mathbb{E}_{(s,p,o) \sim \mathcal{D}} \sum_{t=1}^{|x_{s,p,o}|} \log \pi_{\theta}(x_i | x_{<i}). \quad (5)$$

Table 1: Overall Comparison

Method	PopQA				Wikidata-knowledge infusion			
	ACC (%)	HR (%)	MRR (%)	Prob (%)	ACC (%)	HR (%)	MRR (%)	Prob (%)
Qwen3-4B	13.05	31.60	15.17	12.38	20.65	66.96	31.49	4.44
Prompt (CoT)	10.92	32.47	14.78	10.92	20.02	65.97	30.12	4.24
DPO <i>initialized from Qwen3-4B</i>	9.67	18.06	11.02	33.88	33.84	47.46	37.54	18.42
Continual Training (CT)	29.97	60.36	32.94	20.92	37.42	79.32	45.89	2.57
Iterative RPO	32.50	49.30	25.12	20.77	30.23	72.23	41.63	1.69
PretrainRL(ours)	37.69	66.39	40.51	31.42	46.20	78.98	54.11	7.16
Qwen3-8B	16.32	37.06	18.60	18.60	23.96	70.43	34.81	5.51
Prompt (CoT)	15.54	37.57	17.97	12.14	23.96	70.43	34.81	5.51
DPO <i>initialized from Qwen3-8B</i>	13.08	22.17	13.92	42.11	38.51	49.49	41.70	51.41
Continual Training (CT)	38.53	69.17	41.93	27.50	31.45	78.78	41.47	4.51
Iterative RPO	41.23	48.13	32.28	40.83	33.50	75.13	44.71	1.98
PretrainRL(ours)	48.61	74.92	51.51	40.48	47.58	77.10	55.98	11.10
Llama3-8B	23.66	48.69	26.35	23.10	27.60	74.14	38.43	7.54
Prompt (CoT)	24.28	49.68	27.06	21.32	27.71	73.72	37.77	6.15
DPO <i>initialized from Llama3-8B</i>	15.68	26.02	16.93	42.28	54.42	62.54	57.12	28.39
Continual Training (CT)	13.40	33.18	16.53	20.10	53.16	89.35	63.70	18.33
Iterative RPO	33.83	32.63	22.64	28.17	58.92	79.72	59.24	5.34
PretrainRL(ours)	50.16	51.38	38.63	30.67	64.69	80.79	70.74	11.01

This CT loss serves as a regularization term, anchoring the model’s distribution to the high-quality chosen responses (i.e., ground truth o). The final loss function is a linearly weighted sum of the two components:

$$\mathcal{L}_{\text{PretrainRL}} = \mathcal{L}_{\text{DPO}} + \lambda \mathcal{L}_{\text{CT}} \quad (6)$$

3.4 Evaluation metrics

The previous analysis indicates that the imbalanced training of knowledge results in low conditional probability $P(o|s, p; \theta)$ of tail knowledge. To evaluate the mastery of knowledge, we can ask LLM to complete a continuation or QA tasks given context s, p . However, traditional metrics, such as accuracy, cannot intuitively evaluate the overall state of probabilities. Therefore, we apply the **beam search** decoding method, aligning with proposed negative sampling, and introduce three metrics to comprehensively address this challenge. Suppose the beam size is k , we define the metrics as follows.

- **Hit Ratio (HR@k):** HR is a binary metric that indicates whether LLMs output at least one correct answer among all beam search responses. It demonstrates the diversity, which is commonly high for a base model.
- **Mean Reciprocal Rank (MRR@k):** MRR evaluates a beam search list by averaging the reciprocal of the rank (ordered by Prob) of the first correct answer across all responses. This is a metric that balances between Prob and HR.

- **Probability (Prob@k)** This value reflects the model’s confidence in a particular continuation. It only counts when the responses of beam search hit the correct answer.

Equipped with these metrics (i.e., Acc, HR, MRR, Prob), we can comprehensively assess the LLMs’ mastery of knowledge. Furthermore, such metrics are also friendly to base model, which is beneficial for exploration of hallucination-mitigation methods in pretraining phase.

We demonstrate Qwen’s mastery of knowledge on POPQA, using the above metrics. The results show that the metric HR is high, indicating that Qwen has trained this knowledge. Nevertheless, the metrics MRR and Prob are relatively low, leading to a low accuracy of this knowledge (i.e., factual hallucination). This phenomenon is much more severe for tail knowledge, according to the previous discussion of tail and head knowledge. Moreover, scaling up the model size has limited effect on mitigating this issue.

4 Experiments

4.1 Experimental Setup

Datasets We use POPQA (Mallen et al., 2023), Wikidata-knowledge infusion (Lv et al., 2025), and EntityQuestions (Sciavolino et al., 2022) as the datasets in our experiments. These datasets primarily focus on evaluating the factual knowledge of models, and all of them exhibit long-tail distributions. More detailed datasets are provided in Appendix B.1.

Models We select the Qwen3-4B, Qwen3-7B, Qwen3-14B (Yang et al., 2025) and Llama3-8B (Grattafiori and et al., 2024) for our experiments. In our experiments, we utilize the base versions of all models, which have not undergone instruction tuning or other forms of post-training. As discussed earlier, our method focuses on consolidating factual knowledge during the pretraining stage while minimizing impacts on other knowledge. In contrast, instruct-tuned models have undergone extensive post-training, which make their embedded knowledge more rigid and difficult to modify. Besides, fine-tuning instruct model on such tasks often yields marginal improvements and can sometimes lead to performance degradation (Wu et al., 2025). We apply our method to both base and instruct-tuned models. The results in Table 2 demonstrate that Base-version models provide greater "learnable room" for improvement compared to instruct models when applying our approach.

Table 2: Compare Base and Instruct Model

Method	ACC (%)	HR (%)	MRR (%)	Prob (%)
Qwen3-4B-Instruct	10.70	24.09	12.32	32.30
w/ PretrainRL	33.21	60.62	35.93	36.90
Qwen3-4B-Base	13.05	31.60	15.17	12.38
w/ PretrainRL	37.69	66.39	40.51	31.42
Qwen3-4B-Instruct	18.09	55.17	26.44	20.83
w/ PretrainRL	41.59	73.62	50.94	24.21
Qwen3-4B-Base	20.65	66.96	31.49	4.44
w/ PretrainRL	46.20	78.98	54.11	7.16

Implementation Details In our method, we first evaluate the base model on the POPQA (Mallen et al., 2023), Wikidata-Knowledge Infusion (Lv et al., 2025), and EntityQuestions (Sciavolino et al., 2022) datasets, while simultaneously recording the model’s own knowledge recall for these questions. Following the procedure described in Section 3, we then compute the high-frequency attribute values the model tends to produce for this class of questions and use them as the pool of negative candidates. In all experiments, for each training instance we randomly sample **five negatives** from the top 20 candidate pool, yielding five positive–negative pairs for each question. We repeat the POPQA dataset three times during training, expanding it from 13k to 38k samples. Similarly, the EntityQuestions dataset is repeated twice, increasing its size from 17k to 34k samples. All experiments are conducted for one epoch of training. We use a simple template “{Question}{pos}<pad>{Question}{neg}” to format all of preference pairs (more details can be seen Ap-

pendix A). We employ the AdamW optimizer with a fixed weight decay of 0.1. The learning rate is set to $3e-5$ for POPQA and Wikidata-Knowledge Infusion, and to $2e-5$ for EntityQuestions. The coefficient β in the DPO loss is tuned in 0.05, 0.1, 0.5, 1.0, and we end up using 0.1 in this experiment. For a fair comparison, all experiments share identical training hyperparameters. After training, we evaluate model performance using ACC, Hit, MRR, and Prob with beam size $k = 50$, as defined in Section 3.4. We evaluate all models with the temperature set to 0. All training was conducted on eight H800 GPUs.

4.2 Main Results

PretrainRL improves over baselines. We compare against Chain-of-Thought(CoT) (Wei et al., 2023), Direct Preference Optimization(DPO) (Rafailov et al., 2023), Continued Training(CT), and iterative RPO (Pang et al., 2024). We apply the standard DPO to the same set of preference pairs, initializing from the Qwen3-4B-Base. We also evaluate several mainstream closed-source LLMs, which have larger parameter scales and exhibit advanced reasoning capabilities. As shown in Figure 1, despite their impressive capabilities, these closed-source LLMs perform poorly on long-tail datasets like POPQA, exhibiting high hallucination rates (Sun et al., 2024). In Table 1, we demonstrate that PretrainRL substantially enhances the model’s factual memorization and significantly outperforms baselines as well as the evaluated closed-source LLMs. Notably, the improvements are even more pronounced for long-tail distribution dataset.

Table 3: Scaling model size

Method	ACC (%)	HR (%)	MRR (%)	Prob (%)
Qwen3-4B	13.05	31.60	15.17	12.38
w/ PretrainRL	37.69	66.39	40.51	31.42
Qwen3-8B	16.32	37.06	18.60	18.60
w/ PretrainRL	48.61	74.92	51.51	40.48
Qwen3-14B	19.41	41.57	21.90	20.10
w/ PretrainRL	63.44	84.28	65.58	48.08

4.3 Scalability Study

To further validate the effectiveness of our method, we conduct experiments on a larger-parameter model (Qwen3-14B-Base). As shown in Table 3, PretrainRL consistently enhances performance even for models with substantially larger parameter. The increased parameter inherently strengthens the model’s learning capability and provides greater room for knowledge consolidation. This enables

PretrainRL to more effectively unlock and refine the latent knowledge embedded in the LLM, particularly leveraging the enhanced representational power of larger architectures.

Table 4: PretrainRL on EntityQuestions

Method	ACC (%)	HR (%)	MRR (%)	Prob (%)
Qwen3-4B	16.39	44.53	20.49	7.88
w/ PretrainRL	28.22	56.23	31.48	23.17
Qwen3-8B	21.45	51.97	25.66	8.90
w/ PretrainRL	31.39	61.87	35.02	25.46

To further validate the effectiveness of our method on a larger-scale dataset, we evaluate it on EntityQuestions, a dataset sharing a long-tail distribution with POPQA but containing over 1.75 million questions. For each instance, we generate five negative example pairs, resulting in 8.7 million training instances for PretrainRL. As shown in Table 4, both models exhibit significant improvements on EntityQuestions, aligning with our results on other datasets. This consistency underscores the scalability and generalizability of our approach across varying model sizes and data scales.

4.4 Ablation Study

In this section, we conduct ablation studies to investigate the contributions of each component of PretrainRL to model performance. In our experiments, w/o NTP removes the next-token prediction (NTP) loss term from the total loss and computes only the DPO loss on the preferred pairs for preference optimization; w/o DPO removes the DPO loss and computes only the NTP loss, equivalent to standard continual pretraining. For fairness, all experiments use the same amount data. As reported in Table 5, removing the NTP loss (w/o NTP) reduces accuracy and hit rate but substantially increasing the probability of generating chosen (winning) responses. Meanwhile, removing the DPO loss (w/o DPO) helps consolidate the model’s factual knowledge, but in essence increases the frequency of long-tail knowledge by optimizing the corpus distribution. However, this approach risks suppressing other knowledge, as evidenced by the decreased Prob metric on the Wikidata-knowledge infusion dataset. In contrast, PretrainRL balances these trade-offs by learning knowledge during continual pretraining while reshaping the model’s probability distribution to make room for learning truthful responses, ultimately yielding comprehensive improvements.

We compare different sampling strategies to val-

Table 5: Ablation Study

Method	ACC (%)	HR (%)	MRR (%)	Prob (%)
<i>PopQA</i>				
Qwen3-4B	13.05	31.60	15.17	12.38
w/o NTP	9.67	18.06	11.02	33.88
w/o DPO	29.97	60.36	32.94	20.92
w/ PretrainRL	37.69	66.39	40.51	31.42
Qwen3-8B	16.32	37.06	18.60	18.60
w/o NTP	13.08	22.17	13.92	42.11
w/o DPO	38.53	69.17	41.93	27.50
w/ PretrainRL	48.61	74.92	51.51	40.48
<i>wikidata</i>				
Qwen3-4B	20.65	66.96	31.49	4.44
w/o NTP	33.84	47.46	37.54	18.42
w/o DPO	37.42	79.32	45.89	2.57
w/ PretrainRL	46.20	78.98	54.11	7.16
Qwen3-8B	23.96	70.43	34.81	5.51
w/o NTP	38.51	49.49	41.70	51.41
w/o DPO	31.45	78.78	41.47	4.51
w/ PretrainRL	47.58	77.10	55.98	11.10

idate the effectiveness of our rejected sampling method. Concretely, we partition POPQA by quantiles of attribute-value popularity, designating the top 10% of attribute values by popularity as *head* attributes and treating the remainder as *mid-tail* attributes. For each instance, rejected samples are randomly drawn from the head attribute pool. Aside from the sampling strategy, all experimental settings are kept identical. As shown in Appendix C, our method demonstrates superior performance. Popularity-based sampling relies on an externally defined popularity metric and a rule-based partitioning of high-frequency knowledge, making it susceptible to dataset-specific distribution biases. In contrast, PretrainRL identifies high-frequency knowledge based on the model’s intrinsic probability distribution during generation, more faithfully reflecting the model’s internal knowledge.

Table 6: Performance on different benchmarks

Method	CEval	MMLU	MATH	GSM8K	BBH
Qwen3-4B	50.81	52.39	14.46	83.62	71.80
w/ PretrainRL	49.06	52.01	17.22	80.97	66.07
w/ SFT	29.38	31.76	0.26	0.53	4.02

4.5 Generalization Study

Prior studies (Charton and Kempe, 2024; Lv et al., 2025) have shown that excessive knowledge infusion does not always yield proportional benefits, and fine-tuning models for specific capabilities may inadvertently compromise performance on downstream tasks. To validate the generalizability and robustness of our method, we evaluate model trained with PretrainRL on diverse benchmarks: CEval (Huang et al., 2023), MMLU (Hendrycks et al., 2021), MATH (Wang et al., 2024b), GSM8K

(Cobbe et al., 2021), and BBH (Suzgun et al., 2022). All evaluations are conducted using the OpenCompass (Contributors, 2023). As shown in Table 6, PretrainRL does not cause any noticeable decline in overall performance, indicating that the improvements are not achieved at the expense of other abilities. In contrast, models fine-tuned via SFT exhibit substantial performance collapse on downstream tasks, highlighting PretrainRL’s unique ability to balance knowledge consolidation with general capability preservation.

4.6 Case Study

The imbalanced data distribution in training corpora leads LLM to exhibit a strong bias toward generating head knowledge when answering questions, while long-tail knowledge often degrades from conditional probability estimation to marginal probability approximation, resulting in severe hallucinations. In our method, we employ beam search decoding to obtain negative samples while simultaneously capturing the probability distribution of candidate answers. For instance, for the question “What is Yungay the capital of?”, beam search retrieves the top 50 candidate answers, including the ground truth “Yungay Province”, which ranks 29th with a probability of $5.5802e - 33$. This indicates that while the model retains latent memory of the correct answer, its marginal probability estimation fails to prioritize truthful responses. PretrainRL addresses this by debiasing the model’s probability distribution for long-tail knowledge, thereby mitigating hallucinations. Figure 2 visualizes the probability distribution before and after training for a representative instance. At the beginning, the model initially assigns negligible probability to the correct answer. After PretrainRL, the answer’s probability is significantly recalibrated and stabilized, demonstrating the method’s effectiveness in aligning knowledge recall with factual accuracy.

5 Related work

Hallucination has been a fundamental challenge in LLMs and has received extensive attention in existing studies (Sahoo et al., 2024; Kalai et al., 2025; Liu et al., 2025). At the data level, prior research (Kamalloo et al., 2023; Pal et al., 2023) seek to reduce the interference caused by noise, factual inaccuracies, and conflicting data by constructing higher-quality training corpora. Nevertheless, corpus imbalance is difficult to avoid. At the training

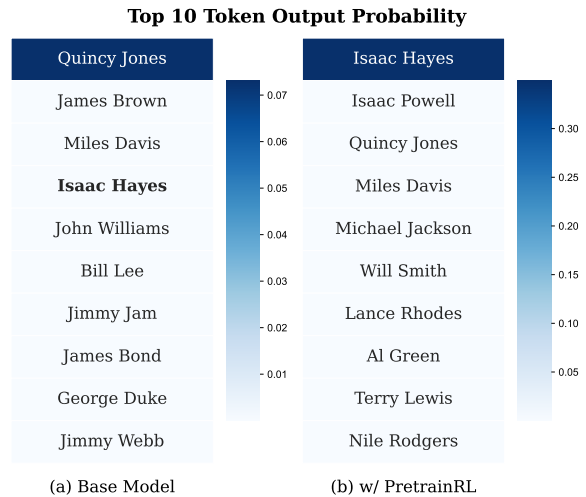


Figure 2: The output of the model for the question “Who was the composer of Shaft?”.

strategy level, most work (Song et al., 2025) focuses on enabling models to say “I don’t know” when they lack confidence in answering a question. Along this line, several approaches employ self-refinement mechanisms where models improve output accuracy through post-hoc feedback or internal reasoning processes. In addition, knowledge injection and Reinforcement Finetuning approaches have been explored to mitigate hallucinations. Notably, recent investigations (Gudibande et al., 2023; Gekhman et al., 2024) reveal that fine-tuning on unpopular factual knowledge may encourage hallucinations and LLMs struggle to integrate new knowledge through fine-tuning. To alleviate hallucinations concerning unpopular knowledge, we conduct an in-depth investigation from the perspective of the pretraining paradigm.

6 Conclusion

In this paper, we identify the imbalanced data distribution in pretraining corpora as a root cause of factual hallucinations. To address this at its source, we propose PretrainRL, a novel framework that integrates reinforcement learning into the pretraining phase. Guided by a “debiasing then learning” principle, PretrainRL leverages DPO and an efficient negative sampling strategy to reshape the model’s probability distribution, making room for factual knowledge to be consolidated. Extensive experiments on three public benchmarks demonstrate that PretrainRL significantly alleviates factual hallucinations. Our work presents a paradigm shift towards a more fundamental “pretrain-and-align” approach for instilling factual reliability, moving beyond post-hoc corrections.

Limitations

In our sampling procedure, to improve efficiency, we aggregate the head knowledge generated by the model for each question category and select the corresponding head knowledge as negative samples, which presupposes that the dataset offers a meaningful distinction among question categories. If the dataset lacks such differentiation, the sampling strategy can be flexibly substituted with alternative methods such as RLVR to maintain robustness. This adaptability ensures that PretrainRL remains applicable across diverse data structures while preserving its core objective of mitigating hallucinations through targeted negative sampling.

References

- Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, and 1 others. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*.
- Lang Cao. 2023. Learn to refuse: Making large language models more controllable and reliable through knowledge scope limitation and refusal mechanism. *arXiv preprint arXiv:2311.01041*.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2021. [Editing factual knowledge in language models](#). *Preprint*, arXiv:2104.08164.
- François Charton and Julia Kempe. 2024. [Emergent properties with repeated examples](#). *Preprint*, arXiv:2410.07041.
- Zhe Chen, Yuehan Wang, Bin Zhao, Jing Cheng, Xin Zhao, and Zongtao Duan. 2020. Knowledge graph completion: A review. *Ieee Access*, 8:192435–192456.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- OpenCompass Contributors. 2023. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>.
- DeepSeek-AI. 2024. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Qingxiu Dong, Li Dong, Yao Tang, Tianzhu Ye, Yutao Sun, Zhifang Sui, and Furu Wei. 2025. Reinforcement pre-training. *arXiv preprint arXiv:2506.08007*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, and 1 others. 2020. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig. 2024. [Does fine-tuning LLMs on new knowledge encourage hallucinations?](#) In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7765–7784, Miami, Florida, USA. Association for Computational Linguistics.
- Aaron Grattafiori and Abhimanyu Dubey et al. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Arnav Gudibande, Eric Wallace, Charlie Snell, Xinyang Geng, Hao Liu, Pieter Abbeel, Sergey Levine, and Dawn Song. 2023. [The false promise of imitating proprietary llms](#). *Preprint*, arXiv:2305.15717.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, and 1 others. 2021. Knowledge graphs. *ACM Computing Surveys (Csur)*, 54(4):1–37.
- Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. 2023. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. In *Advances in Neural Information Processing Systems*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM computing surveys*, 55(12):1–38.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, and 1 others. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Adam Tauman Kalai, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang. 2025. [Why language models hallucinate](#). *Preprint*, arXiv:2509.04664.

- Ehsan Kamaloo, Aref Jafari, Xinyu Zhang, Nandan Thakur, and Jimmy Lin. 2023. [Hagrid: A human-llm collaborative dataset for generative information-seeking with attribution](#). *Preprint*, arXiv:2307.16883.
- Zixuan Ke, Yijia Shao, Haowei Lin, Tatsuya Konishi, Gyuhak Kim, and Bing Liu. 2023. Continual pre-training of language models. *arXiv preprint arXiv:2302.03241*.
- Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, and 1 others. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Teaching models to express their uncertainty in words. *arXiv preprint arXiv:2205.14334*.
- Langming Liu, Haibin Chen, Yuhao Wang, Yujin Yuan, Shilei Liu, Wenbo Su, Xiangyu Zhao, and Bo Zheng. 2025. Eckgbench: Benchmarking large language models in e-commerce leveraging knowledge graph. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pages 6461–6465.
- Kangtao Lv, Haibin Chen, Yujin Yuan, Langming Liu, Shilei Liu, Yongwei Wang, Wenbo Su, and Bo Zheng. 2025. [How to inject knowledge efficiently? knowledge infusion scaling law for pre-training large language models](#). *Preprint*, arXiv:2509.19371.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. *arXiv preprint arXiv:2303.08896*.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022a. Locating and editing factual associations in gpt. *Advances in neural information processing systems*, 35:17359–17372.
- Kevin Meng, Arnab Sen Sharma, Alex Andonian, Yonatan Belinkov, and David Bau. 2022b. Mass-editing memory in a transformer. *arXiv preprint arXiv:2210.07229*.
- Eric Mitchell, Charles Lin, Antoine Bosselut, Chelsea Finn, and Christopher D Manning. 2021. Fast model editing at scale. *arXiv preprint arXiv:2110.11309*.
- OpenAI. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. 2023. [Med-halt: Medical domain hallucination test for large language models](#). *Preprint*, arXiv:2307.15343.
- Richard Yuanzhe Pang, Weizhe Yuan, He He, Kyunghyun Cho, Sainbayar Sukhbaatar, and Jason E Weston. 2024. [Iterative reasoning preference optimization](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, and 1 others. 2018. Improving language understanding by generative pre-training.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? *arXiv preprint arXiv:2002.08910*.
- Pranab Sahoo, Prabhash Meharia, Akash Ghosh, Sriparna Saha, Vinija Jain, and Aman Chadha. 2024. [A comprehensive survey of hallucination in large language, image, video and audio foundation models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 11709–11724, Miami, Florida, USA. Association for Computational Linguistics.
- Christopher Sciavolino, Zexuan Zhong, Jinhyuk Lee, and Danqi Chen. 2022. [Simple entity-centric questions challenge dense retrievers](#). *Preprint*, arXiv:2109.08535.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Haizhou Shi, Zihao Xu, Hengyi Wang, Weiyi Qin, Wenjuan Wang, Yibin Wang, Zifeng Wang, Sayna Ebrahimi, and Hao Wang. 2024. Continual learning of large language models: A comprehensive survey. *ACM Computing Surveys*.
- Linxin Song, Taiwei Shi, and Jieyu Zhao. 2025. [The hallucination tax of reinforcement finetuning](#). *Preprint*, arXiv:2505.13988.
- Kai Sun, Yifan Ethan Xu, Hanwen Zha, Yue Liu, and Xin Luna Dong. 2024. [Head-to-tail: How knowledgeable are large language models \(llms\)? a.k.a. will llms replace knowledge graphs?](#) *Preprint*, arXiv:2308.10168.

- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, , and Jason Wei. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- SMTI Tonmoy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. 2024. A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313*, 6.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Liyuan Wang, Xingxing Zhang, Hang Su, and Jun Zhu. 2024a. A comprehensive survey of continual learning: Theory, method and application. *IEEE transactions on pattern analysis and machine intelligence*, 46(8):5362–5383.
- Peiyi Wang, Lei Li, Zhihong Shao, Runxin Xu, Damai Dai, Yifei Li, Deli Chen, Yu Wu, and Zhifang Sui. 2024b. [Math-shepherd: Verify and reinforce LLMs step-by-step without human annotations](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9426–9439, Bangkok, Thailand. Association for Computational Linguistics.
- Song Wang, Yaochen Zhu, Haochen Liu, Zaiyi Zheng, Chen Chen, and Jundong Li. 2024c. Knowledge editing for large language models: A survey. *ACM Computing Surveys*, 57(3):1–37.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Taiqiang Wu, Runming Yang, Jiayi Li, Pengfei Hu, Yik-Chung Wu, Ngai Wong, and Yujiu Yang. 2025. [Shadow-ft: Tuning instruct model via training on paired base model](#). *Preprint*, arXiv:2505.12716.
- Ziwei Xu, Sanjay Jain, and Mohan Kankanhalli. 2024. Hallucination is inevitable: An innate limitation of large language models. *arXiv preprint arXiv:2401.11817*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, and 1 others. 2025. Siren’s song in the ai ocean: A survey on hallucination in large language models. *Computational Linguistics*, pages 1–46.
- Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong Wu, Jingjing Xu, and Baobao Chang. 2023. Can we edit factual knowledge by in-context learning? *arXiv preprint arXiv:2305.12740*.

A Prompt

A.1 Question Answering Prompt (Few-shot)

Table 7: Question answering (Few-shot).

Answer the following questions in as few words as possible.

Question: What is the capital of China?
Answer: Beijing

Question: What’s the job of Song Kang-ho in Parasite (2019)?
Answer: actor

Question: {QUESTION}
Answer:

A.2 Question Answering Prompt (CoT)

Table 8: Question answering (CoT).

Your task is to answer the question below. Give step by step reasoning before you answer, and when you’re ready to answer. Answer the question in as few words as possible.

Question: What is the capital of China?
Answer: Beijing

Question: What’s the job of Song Kang-ho in Parasite (2019)?
Answer: actor

Question: {QUESTION}
Answer:

A.3 Training data template

Table 9: Training data template

{Question}{pos}<pad>{Question}{neg}

Example:
What is the capital of Arno? Florence<pad>What is the capital of Arno?
Rome

B Dataset

B.1 Dataset Description

POPQA¹ is created by sampling knowledge triples from Wikidata and converting them to natural language questions, followed by popularity calculation. The whole dataset consists of 14k questions. To avoid knowledge conflicts during the learning process, we deduplicated the POPQA dataset by retaining only subject entities that have a unique corresponding property entity, resulting in a final set of 13k questions. Wikidata-knowledge infusion² follows a similar construction process to POPQA, aggregating factual triplets from Wikidata and selecting six common relationship types to form 28k

¹<https://github.com/AlexTMallen/adaptive-retrieval>

²https://huggingface.co/datasets/RJZ/wikidata_triple_en

questions. To balance category disparities, we perform sampling on each category, yielding a final dataset of 16k questions. To validate the effectiveness of our method on a larger-scale, we also conduct experiments on EntityQuestions³, another popular open-domain QA dataset. Similar to POPQA, EntityQuestions uses Wikipedia hyperlink counts as a proxy for entity frequency and samples knowledge triples from Wikidata based on frequency distributions. After deduplicating subject entities, we obtain a final set of 1.75 million questions. The aforementioned datasets vary in size but all feature long-tail distributions. Following in POPQA, we analyze the popularity distribution of attribute values.

B.2 Statistics of the evaluation dataset

Table 10: Statistics of POPQA

Relation	Question template	Count
occupation	What is {s}’s occupation?	531
father	Who is the father of {s}?	562
country	In what country is {s}?	824
director	Who was the director of {s}?	1784
composer	Who was the composer of {s}?	880
place of birth	In what city was {s} born?	584
color	What color is {s}?	34
religion	What is the religion of {s}?	326
capital of	What is {s} the capital of?	255
producer	Who was the producer of {s}?	1320
author	Who is the author of {s}?	1408
genre	What genre is {s}?	1530
screenwriter	Who was the screenwriter for {s}?	1671
mother	Who is the mother of {s}?	187
capital	What is the capital of {s}?	622
Total-Count		12,518

Table 11: Statistics of Wikidata-knowledge infusion

Relation	Question template	Count
capital	What is the capital of {s}?	421
industry	What is the industry of {s}?	4000
location	What is the location of {s}?	4000
material	What is the material of {s}?	4000
color	What is the color of {s}?	3627
shape	What is the shape of {s}?	342
Total-Count		16,390

³<https://github.com/princeton-nlp/EntityQuestions>

Table 12: Statistics of EntityQuestions

Relation	Question template	Count
P36	What is the capital of {s}?	6704
P407	Which language was {s} written in?	5113
P26	Who is {s} married to?	7979
P159	Where is the headquarter of {s}?	7994
P276	Where is {s} located?	7902
P40	Who is {s}'s child?	7981
P176	Which company is {s} produced by?	7913
P20	Where did {s} die?	7985
P112	Who founded {s}?	4060
P127	Who owns {s}?	7936
P19	Where was {s} born?	7991
P740	Where was {s} founded?	7506
P413	What is {s} famous for?	7971
P800	What position does {s} play?	1762
P69	Where was {s} educated?	7990
P50	Who is the author of {s}?	7945
P170	Who was {s} created by?	6709
P106	What kind of work does {s} do?	7988
P131	Where is {s} located?	7921
P17	Which country is {s} located in?	7976
P175	Who performed {s}?	7767
P136	What type of music does {s} play?	7954
P264	What music label is {s} represented by?	7841
P495	Which country was {s} created in?	7909
Total-Count		174,797

C Compare different sampling methods

Table 13: Compare sampling methods on POPQA

Method	ACC (%)	HR (%)	MRR (%)	Prob (%)
Qwen3-4B	16.39	44.53	20.49	7.88
Popularity-based sampling	24.75	48.98	26.93	27.72
PretrainRL(ours)	37.69	66.39	40.51	31.42
Qwen3-8B	21.45	51.97	25.66	8.90
Popularity-based sampling	35.47	60.97	37.93	34.95
PretrainRL(ours)	48.61	74.92	51.51	40.48