

MaFeRw: Query Rewriting with Multi-Aspect Feedbacks for Retrieval-Augmented Large Language Models

Yujing Wang^{1,2}, Hainan Zhang^{1,2*}, Liang Pang⁴, Binghui Guo²,
Hongwei Zheng³, Zhiming Zheng^{1,2}

¹Beijing Advanced Innovation Center for Future Blockchain and Privacy Computing

²School of Artificial Intelligence, Beihang University, China

³Beijing Academy of Blockchain and Edge Computing, China

⁴Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China
{wangyujing, zhanghainan}@buaa.edu.cn

Abstract

In a real-world RAG system, the current query often involves spoken ellipses and ambiguous references from dialogue contexts, necessitating query rewriting to better describe user's information needs. However, traditional context-based rewriting has minimal enhancement on downstream generation tasks due to the lengthy process from query rewriting to response generation. Some researchers try to utilize reinforcement learning with generation feedback to assist the rewriter, but this sparse rewards provide little guidance in most cases, leading to unstable training and generation results. We find that user's needs are also reflected in the gold document, retrieved documents and ground truth. Therefore, by feeding back these multi-aspect dense rewards to query rewriting, more stable and satisfactory responses can be achieved. In this paper, we propose a novel query rewriting method MaFeRw, which improves RAG performance by integrating multi-aspect feedback from both the retrieval process and generated results. Specifically, we first use manual data to train a T5 model for the rewriter initialization. Next, we design three metrics as reinforcement learning feedback: the similarity between the rewritten query and the gold document, the ranking metrics, and ROUGE between the generation and the ground truth. Inspired by RLAIF, we train three kinds of reward models for the above metrics to achieve more efficient training. Finally, we combine the scores of these reward models as feedback, and use PPO algorithm to explore the optimal query rewriting strategy. Experimental results on two conversational RAG datasets demonstrate that MaFeRw achieves superior generation metrics and more stable training compared to baselines.

Code — <https://github.com/TAP-LLM/MaFeRw>

Introduction

Retrieval-augmented generation (RAG) effectively addresses the issue of factual inaccuracy in large language models (LLMs) by integrating relevant retrieved information. It is widely used in various fields such as QA systems (Du and Ji 2022; Siriwardhana et al. 2023; Sharma et al. 2024), content generation (Zhao et al. 2024; Su et al.

* Corresponding author.

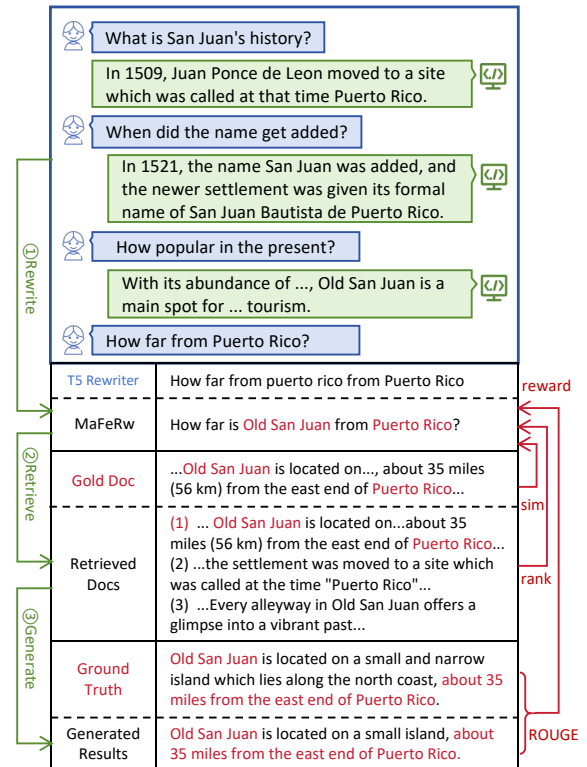


Figure 1: An example of MaFeRw serving RAG and the comparison with T5 rewriter. Green lines represent the inference process of rewriter and RAG, while Red lines indicate the three types of reward metrics feedback to MaFeRw.

2024; Wang et al. 2023), and virtual agents (Yang, Yue, and He 2023; Schick et al. 2023; Zhang 2023). Real-world RAG systems often rely on multi-turn dialogues, where queries include spoken ellipses and ambiguous references from dialogue contexts (Mao et al. 2023b), making it challenging for RAG systems to accurately understand user intent (Ye et al. 2024). In light of this challenge, effective query rewriting is necessary to better describe user information needs and ensure the accuracy of retrieval and generation in RAG.

Currently, RAG mainly performs context-based query

rewriting to improve the quality of response generation, but the resulting enhancement is minimal due to the lengthy process from query rewriting to response generation. Although context-based query rewriting can directly affect the retrieval results, this impact is difficult to transmit to the final generated results through this loose-coupling retrieval and generation framework of RAG. As shown in Figure 1, the T5 rewriter (Lin et al. 2020) trained on context-based rewrite dataset fails to effectively capture the user’s intent. To address this issue, Ma et al. try to use reinforcement learning (RL) with generation feedback to assist the rewriter. They optimize the rewriter with policy gradient and use the accuracy and ‘hit’ rate of RAG inference’s answers to form the reward. However, this reward is sparse and provides little guidance in most cases, resulting in unstable training and generation results of their model. Additionally, it is inefficient to directly use the full RAG inference to obtain RL feedback. Because RL typically requires a large number of training steps to achieve good results, and each inference in RAG takes a long time. Thus, if a stable and efficient method for training the rewriter can be designed, it will undoubtedly further enhance the effectiveness of RAG.

We find that user information needs are also reflected in the gold document, retrieved documents and ground truth, which allows for the design of multi-aspect rewards to provide more dense feedback signals. As illustrated in Figure 1, the current query ‘How far from Puerto Rico’ omits ‘Old San Juan’ from the previous context. The user intent ‘Old San Juan’s distance from Puerto Rico’ is evident not only in the context but also in the gold document, retrieved documents, and ground truth. If query rewriting can leverage this information, it can better assist the model in generating correct answers. Therefore, by feeding multi-aspect dense rewards from the gold document, ranking measures, and generation metrics back to rewriting, more stable and satisfactory responses can be achieved.

In this paper, we propose MaFeRw, a novel query Rewriting method, to improve RAG performance by integrating Multi-aspect Feedback from both the retrieval process and generated results. Specifically, we first use manual data to train a T5 model (Raffel et al. 2020) as the rewriter’s initialization. Next, we design three metrics as reinforcement learning feedback: the similarity between the rewritten query and the gold document, the ranking metrics of similarity between retrieved documents and the ground truth, and ROUGE between the generated response and the ground truth. Inspired by RLAIF (Lee et al. 2023), to achieve more efficient training, we collect datasets based on these metrics and train three kinds of reward models for them. Additionally, we use ROUGE scores of model-rewritten queries and manual-rewritten queries as the fourth feedback to measure the rewriter’s performance. Finally, we combine scores of three reward models and the rewritten ROUGE as feedback, and use PPO algorithm (Schulman et al. 2017; Ziegler et al. 2020) to explore the optimal query rewriting strategy.

Experimental results on two conversational RAG datasets demonstrate that MaFeRw achieves superior generation metrics compared to baselines. Further analysis shows that multi-aspect dense rewards provide a more stable training

process and generation results than single reward, validating the stability and transferability of MaFeRw.

The innovations of this paper are as follows:

- We find that dense reward feedback can lead to more stable and satisfactory generation results than single reward, and design four rewards from multiple aspects, such as gold document, retrieved documents and ground truth.
- To achieve more efficient RL training, we collect the feedback signals and train corresponding reward models, rather than using full RAG inference like the baseline.
- Experimental results on two RAG datasets demonstrate that MaFeRw achieves superior generation metrics and more stable training than baselines.

Related Work

In dialogue systems, user utterances often contain omissions and ambiguous references (Fang et al. 2022; Zhang et al. 2019). Therefore, a rewriting model is needed to resolve these ambiguities in the current query and recover missing elements (e.g., anaphora) from the context to better articulate the user’s information needs. Current research on query rewriting mainly focuses on conversational search tasks, conversational QA tasks, and RAG tasks.

In conversational search tasks, researchers try to explore effective query rewriting methods to enhance the retrieval process. Some studies focus on extracting user intents from the dialogue context. For example, Lin et al. use human rewrites as labels to train a sequence-to-sequence model as the rewriter. Qian and Dou propose query rewriting explicitly highlights relevant terms in the query context, facilitating contextual modeling and thereby improving the quality of the learned contextual query embeddings. In (Ye et al. 2023), LLM is used as the rewriter and is prompted to complete a “rewrite-then-edit” process based on the conversation history. Others incorporate retrieval feedback to enhance rewriting for retrieval tasks. Mo et al. and Mao et al. both take the feature similarity between the query and the target passage into the optimization objective when training the rewriting model and Wu et al. use the retrieval effect as a reward for reinforcement learning. The rewritten queries are encouraged to achieve better retrieval performance.

In conversational QA tasks, mainstream query rewriting studies (Kim et al. 2021; Chen et al. 2022) focus on feeding the generated results back into the rewriting model to enhance answer accuracy. Kim et al. train the QA model by regularizing the consistency of generated answers so that the answers predicted by the model according to the original questions were similar to the answers generated according to the rewritten questions. Chen et al. use the effect of generated answers as a reward for reinforcement learning, to improve the effect of rewriting on the final generated results.

In conversational RAG task, Ma et al. firstly create a pseudo-dataset for supervised warm-up training, then further optimize the rewriter with a policy gradient RL framework, using the accuracy and ‘hit’ rate of LLMs’ answers to form the reward.

However, for RAG tasks, focusing solely on the retrieval process results in rewrites unsuitable for prompting LLMs

during generation. Conversely, considering only the feedback from the final generated results can lead to unstable training of the rewriting model due to the sparse reward signal. Therefore, we propose MaFeRw, which combines multi-aspect reward feedbacks to produce more stable and satisfactory responses.

Problem Definition

In RAG, each dialogue history h^k contains a sequence of (query, answer) pairs $h^k = (q^1, a^1, \dots, q^{k-1}, a^{k-1})$, where q^i and a^i denote the query and the system generation of the i -th turn. A conversational query q^i can be elliptical and ambiguous. Given the dialogue history h^k and the current query q^k , a query rewriting model R_θ with parameters θ rewrites q^k to q_R^k :

$$q_R^k = R_\theta(h^k, q^k).$$

The rewritten query q_R^k is then fed into an off-the-shelf RAG system’s retriever to search for relevant documents $D^k = \{d_1^k, \dots, d_l^k\}$, where d_l^k denotes the l -th retrieved document of the k -th turn, l is the number of retrieved documents. The current gold document is represented as d_+^k . Subsequently, the rewritten query q_R^k , the retrieved documents D^k , and the (q^{k-1}, a^{k-1}) pair from the previous turn are combined as input to prompt the LLM to generate more satisfactory results. The (q^{k-1}, a^{k-1}) pair is used to instruct the LLM to mimic the style of the dataset by explicitly prompting the LLM to ’follow the style of the example responses in the context (q^{k-1}, a^{k-1}) ’. In this paper, we aim to train the query rewriting model R_θ to make the rewritten query q_R^k better describe user needs and ensure the accuracy of retrieval and generation in RAG.

Approach

Framework Overview

For MaFeRw, the rewriter is trained to transform the current query q^k into q_R^k utilizing the dialogue history h^k , which enables the RAG system to generate more satisfactory results. Firstly, we use manual rewriting data to train a T5 model as the rewriter’s initialization. Secondly, we design three metrics as RL feedback: the similarity between the rewritten query q_R^k and the gold document d_+^k , the ranking metrics of similarity between retrieved documents D^k and the ground truth G^k , and ROUGE between the generation g^k and the ground truth G^k . Inspired by RLAIIF, we train three kinds of reward models for the above metrics on three corresponding paired data sets, as shown in Figure 2a. Considering the obvious improvement of manual rewriting in the generation, ROUGE between manual rewrites and model rewrites is used as another feedback. Finally, we combine scores of these reward models and the rewritten ROUGE to feedback PPO training, exploring the optimal query rewriting strategy, as shown in Figure 2b.

Rewriter Initialization

To equip the rewriter with fundamental rewriting capabilities, we train the pre-trained T5-base model on datasets containing manually rewritten queries. The query rewriter con-

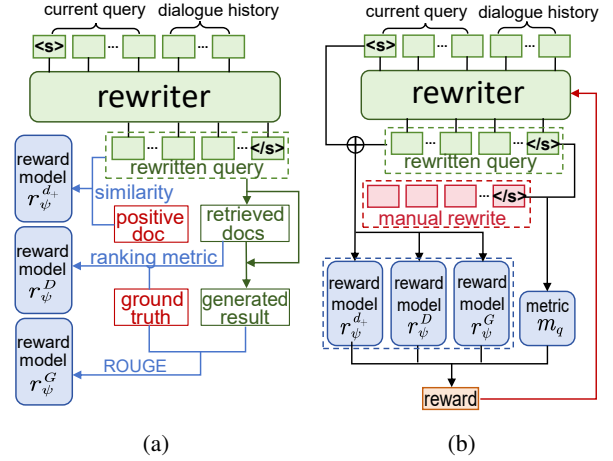


Figure 2: The framework of MaFeRw. (a) Three feedback metrics are: the similarity between rewritten query and gold document, the ranking metric of similarity between ground truth and retrieved documents, and the ROUGE scores between generation and ground truth. Corresponding reward models are trained for these metrics. (b) When training the rewriter using PPO algorithm, the reward is composed of scores from three reward models and the rewritten ROUGE.

catenates the current query q^k with the dialogue history h^k as input. Similar to the approach used by (Wu et al. 2022), a separator token ’[SEP]’ is inserted between each interaction, and dialogue turns are connected in reverse order, as

$$I = [\text{CLS}]q^k[\text{SEP}]a^{k-1}[\text{SEP}]q^{k-1} \dots q^1[\text{SEP}].$$

When q_+ is a manual-rewritten query, the objection of initialization is to optimize parameters θ of the rewriter R_θ by minimizing the cross-entropy loss between the model’s prediction q_R and q_+ :

$$\mathcal{L}_{init} = -y_{q_+} \log y_{q_R},$$

where y_{q_+} is the one-hot vector of q_+ and y_{q_R} is the distribution over tokens in q_R predicted by the rewriter. The rewriter after initialization is represented as R_θ^0 , which defines a probabilistic policy for this task.

Multi-aspect Feedbacks

Since user information needs are also reflected in gold documents, retrieved documents, and ground truth, we design three metrics as RL feedback. For simplicity, we omit the superscript k of q^k , q_R^k , d_+^k , D^k , G^k and g^k in the rest of the paper if not specified.

(1) **The similarity between q_R and d_+** In RAG, the retriever converts the query and documents into dense vectors, calculates similarity scores between the query vector and document vectors, and retrieves the top- l most relevant documents. To ensure the rewritten query q_R retrieves the gold document d_+ more effectively, the cosine similarity (CS) between the dense vector v_{q_R} of q_R and the dense vector v_{d_+} of d_+ is used as a metric m_{d_+} which provides retrieval feedback to the rewriter, i.e., $m_{d_+} = \text{CS}(v_{q_R}, v_{d_+})$.

(2) **The ranking metric of similarity between D and G** Since user needs are embedded in the ground truth G , and the order of retrieved documents D impacts generation quality, we hope the rewritten query can assist the retrieval model in prioritizing the document that exhibits high semantic relevance to the ground truth. Therefore, we propose a ranking metric m_D using G and $D = \{d_1, \dots, d_l\}$. Specifically, we transform the ground truth G into a dense vector v_G using the retrieval embedding model. Then the CS value is calculated between v_G and dense vectors of retrieved documents $\{v_{d_i}\}_{i=1}^l$. For each d_i , let $1/i$ be its rank score. The metric m_D is computed by summing the product of each document’s rank score and its CS value with v_G , formulated as $m_D = \sum_{i=1}^l 1/i \cdot \text{CS}(v_G, v_{d_i})$. A higher m_D indicates that the d_i with greater similarity to G is ranked higher, thereby facilitating the generation of better responses.

(3) **ROUGE between g and G** To ensure that the generated output g produced by RAG closely aligns with the ground truth G , we use ROUGE score between the generation g and the ground truth G as the metric m_G .

(4) **ROUGE between q_R and q_+** Considering the obvious improvement of manual rewrites in the generation, we also use ROUGE between the manual rewrite q_+ and model rewrite q_R as a reference metric m_q .

Reward Model

The end-to-end RL approach without a reward model demands extensive training time due to the prolonged inference time of RAG and the large number of training steps required. Incorporating the complete RAG process into RL led to training durations of several days, severely limiting tuning and flexibility in later stages. To address this issue, reward models were employed to evaluate rewritten queries, effectively decoupling the evaluation from RAG process and enabling more efficient training. Inspired by RLAIIF, we train three types of reward models for the aforementioned metrics on three corresponding paired datasets. The reward model is based on the T5-base model with a value head, possessing significantly fewer parameters than the generation model Llama2-13b. During RL training, the reward is derived from scores of reward models rather than using the complete RAG process, thereby reducing RAG inference time.

Specifically, for the given dialogue history h and the current query q , we pair the corresponding manual rewrite q_m , rewrite q_{T5} generated by the initialized rewriter, and rewrite q_s obtained through sampling. All rewrites are then put into the RAG system to calculate values of m_{d+} , m_D and m_G . For each metric, we compare rewrites in every pair and label the one with the higher metric value as ‘chosen’ and the other as ‘rejected’, creating a paired dataset of (chosen, rejected) rewrites. Taking the paired dataset of metric m_G as an example, if $m_G(q_m) > m_G(q_{T5}) > m_G(q_s)$, then for the dialogue history h and current query q , three pairs of data are collected:

$$\{(chosen : [h + q, q_m], rejected : [h + q, q_{T5}]), \\ (chosen : [h + q, q_m], rejected : [h + q, q_s]), \\ (chosen : [h + q, q_{T5}], rejected : [h + q, q_s])\},$$

here “+” denotes string concatenation. When training the reward model, we adopt a method similar to Wang et al.. By following the Bradley-Terry model (Bradley and Terry 1952), we formulate a preference distribution by employing the reward model r_ψ as outlined below:

$$P_\psi(q_c \succ q_r | h + q) = \sigma(r_\psi(h + q, q_c) - r_\psi(h + q, q_r)), \quad (1)$$

where σ is the logistic function, q_c and q_r represent the chosen and rejected rewrite, respectively. This problem can be treated as a binary classification task, resulting in a negative log-likelihood loss function:

$$\mathcal{L}_{rm} = -\mathbb{E}_{D_{rm}}[\log \sigma(r_\psi(h + q, q_c) - r_\psi(h + q, q_r))], \quad (2)$$

where D_{rm} is the paired dataset.

In this work, we use the initialized rewriter to initialize r_ψ . Additionally, we incorporate an extra linear layer on top of the final transformer layer to generate a scalar prediction representing the reward value. Let r_ψ^{d+} , r_ψ^D and r_ψ^G denote the reward models associated with metrics m_{d+} , m_D and m_G , respectively. Given the dialogue history h , the current query q , and the rewritten query q_R , the corresponding rewards can be obtained as $r_\psi^{d+}(h + q, q_R)$, $r_\psi^D(h + q, q_R)$ and $r_\psi^G(h + q, q_R)$, abbreviated as $r_{d+}(q_R)$, $r_D(q_R)$ and $r_G(q_R)$. To facilitate subsequent aggregation, we scale the scores of reward models to $[0, 1]$.

RL Training

To further train the rewriter, we employ a policy gradient RL framework. In the context of RL, the optimization of the rewriter can be viewed as a Markov Decision Process (MDP) represented by the 5-tuple $\langle S, A, P, r, \gamma \rangle$. Here, the state space S is a finite set constrained by the vocabulary and the sequence length. The action space A comprises the available vocabulary and the transition probability P is determined by the policy network, specifically the rewriter model R_θ^{RL} . The reward function r provides the reward value based on the current state, and γ is the discount factor. For each step t , the current state s^t comprises the dialogue history h , the current query q , and the already generated rewrite $q_{R, [< t]}$. The action a^t is the generation of the next token based on the current state. When the generation process stops, the reward for this episode is calculated using the reward function.

In the RL stage, for the dialogue history h , the current query q , and the rewritten query q_R , we combine scores of these reward models and the metric m_q as the feedback r_{RL} for RL training. This is expressed as:

$$r_{RL}(q_R) = \lambda_1 r_{d+}(q_R) + \lambda_2 r_D(q_R) \\ + \lambda_3 r_G(q_R) + \mu m_q(q_R), \quad (3)$$

where λ_1 , λ_2 and λ_3 are hyperparameters that control the relative importance of each reward model, and μ is a hyperparameter that controls the weight of the metric m_q .

Then we use PPO algorithm to explore the optimal query rewriting strategy R_θ^{RL} . The reward objective we need to maximize in policy optimization is:

$$r_{total} = r_{RL}(q_R) - \eta \text{KL}(R_\theta^{RL} || R_\theta^0), \quad (4)$$

where η is a coefficient governing the magnitude of KL penalty, the initial policy model is the initialized rewriter.

Experiments

Experimental Settings

Datasets We conduct main experiments on two multi-turn dialogue RAG datasets, including QReCC (Anantha et al. 2020) and TopiOCQA (Adlakha et al. 2022). And conduct the transferability experiment on the WSDM@24 Multi-Doc QA dataset¹. QReCC contains 14k conversations with 81k question-answer pairs, built on questions from TREC CAsT (Dalton, Xiong, and Callan 2020), QuAC (Choi et al. 2018), and Google Natural Questions (Kwiatkowski et al. 2019). TopiOCQA is an open-domain conversational dataset with topic switches based on Wikipedia, containing 3,920 conversations with information-seeking questions and free-form answers. WSDM@24 Multi-Doc QA dataset comprises multi-turn question-answering dialogues and corresponding documents from Xiaohongshu.

Evaluation Metrics To evaluate the generated results of RAG, we use four standard evaluation metrics: ROUGE-1, ROUGE-L (Lin 2004), BLEU (Papineni et al. 2002), and METEOR (Banerjee and Lavie 2005). ROUGE prioritizes recall and n-gram overlap, BLEU balances precision and brevity, and METEOR emphasizes semantic similarity. In addition, we utilize MRR (Craswell 2009) to evaluate the retrieval results, as it measures the system’s effectiveness in placing relevant items at the top of the ranked list.

Baselines We compare MaFeRw with four conversational RAG methods: raw RAG (Lewis et al. 2021), T5-base rewriter (Lin et al. 2020), RL-base rewriter (Ma et al. 2023), and ConvGQR (Mo et al. 2023). Raw RAG is the baseline model that performs retrieval and response generation without query rewriting, using the concatenation of dialogue history and the current query. T5-base rewriter uses manual rewrites to train a T5 model for query rewriting based on dialogue history. RL-base rewriter utilizes RL with generation feedback to assist the training of the rewriter. ConvGQR generates potential answers for query expansion and incorporates similarity between the rewrite and gold document into the rewriter’s optimization objective.

Implementation Details In this work, we train the query rewriter based on the pre-trained T5-base model. The dense retrieval component in the RAG system is constructed using FAISS (Johnson, Douze, and Jegou 2021) and the pre-trained embedding model msmarco-roberta-base-ance-firstp (Reimers and Gurevych 2019). The generative model is the Llama-2-13b-chat model (Touvron et al. 2023). After initialization of the rewriter, a value head is added to the rewriter to initialize reward models. Details on hyperparameter determination are provided in the Appendix².

Main Results

We demonstrate experimental results on two datasets.

Accuracy of Reward Models It is crucial to train reward models to accurately score a given rewritten query with dialogue history and correctly categorize it as either “chosen”

¹<https://sites.google.com/view/wsdm24-docqa>

²Appendix can be found at: <https://arxiv.org/abs/2408.17072>

	r_{d+}	r_D	r_G
QReCC	0.8927	0.8534	0.7802
TopiOCQA	0.8362	0.7303	0.7277

Table 1: The accuracy of reward models.

Method	ROUGE-1	ROUGE-L	BLEU	METEOR	MRR
raw RAG	31.98	26.44	5.776	35.10	0.4729
T5-base	35.73	30.50	11.07	38.74	0.4884
RL-base	35.47	29.92	11.27	38.93	0.4742
ConvGQR	34.41	28.33	10.39	37.53	0.4974
MaFeRw	37.05	31.31	12.20	40.76	0.5032
Manual	39.73	34.07	13.54	43.85	0.5649

(a) QReCC

Method	ROUGE-1	ROUGE-L	BLEU	METEOR	MRR
raw RAG	20.76	17.70	2.295	25.48	0.2664
T5-base	21.46	19.12	3.854	25.92	0.3015
RL-base	22.94	20.46	4.993	27.83	0.3717
ConvGQR	17.05	15.11	3.503	22.16	0.3081
MaFeRw	23.97	21.38	5.496	29.51	0.3802
Manual	24.72	21.95	5.367	30.00	0.4001

(b) TopiOCQA

Table 2: The comparison of MaFeRw and the baselines on QReCC and TopiOCQA datasets.

or “rejected”. Table 1 gives information about the accuracy of trained metric-specific reward models on their respective test sets. On both QReCC and TopiOCQA datasets, classification accuracies consistently exceeded 70% for all reward models, ensuring the success of subsequent PPO training based on these reward models.

RAG Performance Table 2 presents a comparison of our method and the baseline approaches in terms of retrieval and generation performance. Compared to the baselines, MaFeRw demonstrates significantly improved retrieval and generation performance on both datasets. MaFeRw achieves 5.07 and 3.19 ROUGE-1 higher than raw RAG on QReCC and TopiOCQA, respectively. Moreover, in both datasets, the improvement in MRR score is accompanied by an improvement in generation metrics.

Compared to QReCC dataset, TopiOCQA dataset features topic shifts within the dialogues. Notably, on TopiOCQA, the rewritten queries show a more significant improvement in MRR compared to using raw RAG with the dialogue history directly for retrieval. MRR of Long Rewriter improves by 6.41% on QReCC compared to raw RAG, and by 42.66% on TopiOCQA. Even T5-base rewriter achieves a more pronounced MRR improvement on TopiOCQA compared to QReCC. However, due to the topic shifts in TopiOCQA, while ConvGQR’s query expansion based on generated answers improves retrieval quality, it fails to effectively enhance generation quality. These findings suggest that query

	ROUGE-1	ROUGE-2	ROUGE-L	MRR
m_{d_+}	36.10	20.70	30.22	0.4863
m_D	35.69	20.40	29.96	0.4924
m_G	36.69	21.59	31.03	0.4937
m_q	36.88	22.06	31.22	0.4897
MaFeRw	37.04	21.84	31.31	0.5032
T5-base rewriter	35.73	20.02	30.50	0.4884

(a) QReCC

	ROUGE-1	ROUGE-2	ROUGE-L	MRR
m_{d_+}	21.27	10.24	18.78	0.3619
m_D	21.86	10.41	19.24	0.3693
m_G	23.14	11.21	20.55	0.3570
m_q	22.92	10.96	20.23	0.3723
MaFeRw	23.97	12.31	21.38	0.3802
T5-base rewriter	21.46	9.874	19.12	0.3015

(b) TopiOCQA

Table 3: The impact of four different metrics on the RL training of the query rewriter.

rewriting can effectively filter out redundant information, thereby significantly enhancing retrieval performance when dealing with dialogues involving topic shifts.

The dialogues in QReCC come from various sources, resulting in more diverse question-and-answer styles than those in TopiOCQA. As a result, RL-base rewriter struggles to adapt to these stylistic variations in QReCC dataset, leading to lower performance metrics compared to T5-base rewriter, and revealing its instability in training and outcomes. These results verify that more satisfactory responses can be obtained by feeding back the results of the remote retrieval and generation modules to query rewriting.

Additionally, it can be observed that the RAG performance with manual rewrites remains consistently strong across these two datasets. This observation underlies our decision to incorporate ROUGE between the rewriter output and the manual rewrite in the RL feedback.

Ablation Study

We conduct RL training on both datasets using feedback from a single reward model or metric and compare results with those of T5-base rewriter. Table 3 demonstrates the impact of four different metrics on rewriter’s RL training.

(1) Impact of the metric m_{d_+} For QReCC, we can see that when only the reward model r_{d_+} is used, ROUGE-1 and ROUGE-2 scores are higher than those of T5-base rewriter but MRR is lower. For TopiOCQA, MRR is much higher than that of T5-base rewriter. This indicates that using m_{d_+} as feedback can, to some extent, guide RL optimization to achieve better retrieval and encourage LLM to improve generation quality, although this promoting effect is not stable.

(2) Impact of the metric m_D Using only the reward model r_D yields a better MRR than T5-base rewriter but




Figure 3: The changes on ROUGE-1 and MRR as training iterations increase when rewrites are applied to RAG.

lower ROUGEs on QReCC. For TopiOCQA, both ROUGE and MRR of r_D are higher than those of T5-base rewriter. This confirms that m_D helps the retriever more effectively find relevant documents, but this benefit does not translate to the generative model, thus failing to produce responses that better meet user needs.

(3) Impact of the metric m_G For both datasets, when only the reward model r_G is used, ROUGE and MRR scores improve compared to T5-base rewriter. On TopiOCQA, ROUGEs show a clear advantage, even though MRR of r_G is smaller than those of r_{d_+} and r_D . This confirms the positive impact of m_G on RAG quality and highlights the necessity of conjunction with other metrics to achieve better results.

(4) Impact of the metric m_q When only the metric m_q is used, both ROUGEs and MRR scores show improvement over T5-base rewriter. On QReCC dataset, the improvement in ROUGE-1 score is notable. This suggests that when facing multi-source, diverse-style conversational data like QReCC, the metric m_q can guide RL optimization to achieve better RAG performance.

These results indicate that individual metrics can improve certain aspects of RAG performance, but their performance varies across different datasets, which is insufficient to stably support RL training. Therefore, combining multiple metrics is necessary to achieve better overall results.

Analysis

To investigate whether using scores output by reward models and m_q as RL feedback can ensure stable training results, we conduct experiments on QReCC. We use either a single metric or reward model to guide the rewriter’s RL training and test the retrieval and generation performance of rewrites after 1000, 1500, 2000, and 2500 iterations. Figures 3a and 3b show the changes in ROUGE-1 and MRR scores as the number of iterations increases, respectively.

It can be observed that when training the rewriter using feedback from the reward model r_{d_+} or r_D , MRR of the rewritten queries in RAG gradually increases, while ROUGE-1 shows slight decreases or fluctuations. This indicates that these two types of reward models can reliably guide the rewriter to optimize towards better retrieval performance, but are insufficient to consistently constrain the rewriter to improve generation results. When utilizing the

Case 1:
Context (QReCC):
q_1 : What are the 2 most abundant elements in the earth’s crust?
a_1 : Oxygen and Silicon are the most abundant elements in the ...
q_2 : How much is Oxygen?
a_2 : Oxygen makes up 46.6% of the Earth’s crust.
q_3 : How much is Silicon?
a_3 : 27.7% of the earth’s crust is silicon.
q_4 : What’s the third most abundant?
a_4 : Aluminum is the third most abundant element in the earth’s crust.
Current Query: What’s the atomic number?
T5-base: What’s the atomic number for silicon? (MRR: 0.0)
RL-base: What’s the atomic number for silicon? (MRR: 0.0)
MaFeRw: What’s the atomic number for aluminum ? (MRR: 1.0)
Generated Response:
T5-base: The atomic number for silicon is 14.
RL-base: The atomic number for silicon is 14.
MaFeRw: The atomic number for aluminum is 13.
Ground Truth: The atomic number of aluminum is 13.
Case 2:
Context (QReCC):
q_1 : Who won the Tour de France in 1985?
a_1 : Bernard Hinault won the Tour de France in 1985.
q_2 : Who won in 2019?
a_2 : Egan Bernal from Team Ineos
q_3 : Who was the youngest winner?
a_3 : The youngest Tour de France winner was Cornet Henri, who was 19 years old.
Current Query: Which year?
T5-base: Which year did Cornet Henri win the Tour de France in 1985? (MRR: 0.5)
RL-base: Which year did Cornet Henri win the 1985 Tour de France? (MRR: 0.5)
MaFeRw: Which year did Cornet Henri win the Tour de France? (MRR: 0.5)
Generated Response:
T5-base: Cornet Henri did not win the Tour de France in 1985.
RL-base: Cornet Henri did not win the Tour de France in 1985. The Tour de France was first held in 1903.
MaFeRw: Cornet Henri won the Tour de France in 1904 .
Ground Truth: 1904 Tour de France.

Figure 4: The case study comparing with two baselines.

reward model r_G for training feedback, although there are fluctuations in ROUGE-1 and MRR metrics, there is an overall upward trend. This trend validates that the reward model can provide positive guidance to the rewriter. When directly using m_q for feedback, ROUGE-1 exhibits a stable increase while MRR shows some variability. This indicates that the metric can, to some extent, consistently guide the rewriter towards optimizing in a manner beneficial for RAG generation. Furthermore, Figure 3 shows the changes in ROUGE-1 and MRR as iterations increase after training the rewriter with synthetic feedback. Both metrics show a more stable upward trend, further validating MaFeRw’s stability.

Based on the above results, it can be analyzed that, by training on paired datasets, reward models can capture a more essential correspondence between rewriting based on dialogue history and each metric. Therefore, RL training with the feedback given by reward models can obtain more stable and effective results than directly training with the whole process of RAG.

Method	ROUGE-1	ROUGE-L	BLEU	METEOR
raw RAG	38.79	23.95	5.989	24.35
T5-base	39.80	24.74	7.227	25.61
MaFeRw	41.01	25.86	8.226	26.46

Table 4: The generation performance on the WSDM@24 Multi-Doc QA task.

Case Study

We present a case study and a comparison with two other rewrite baselines in Figure 4 to illustrate more intuitively the improvement of MaFeRw in capturing users’ needs for multi-turn dialogue RAG.

For the first case, according to the dialogue history, the real need of the current query is to ask for ‘the atomic number of aluminum’. However, because ‘silicon’ is also mentioned in the dialogue history, the two baselines are disturbed and add ‘silicon’ mistakenly. As a result, the error retrieval and generation are triggered. In the second case, both baselines incorrectly add ‘in 1985’ to rewritten queries. Although the correct document can still be retrieved, the LLM is unable to give a correct reply based on the rewritten query due to the interference of the information ‘1985’.

These results verify that MaFeRw can accurately capture the anaphora relationship between the current query and the dialogue history, especially when other interfering words are mentioned in the dialogue history.

Transferability

To evaluate the transferability of MaFeRw, we test its performance on the WSDM@24 Multi-Doc QA dataset by using the rewriting model trained on QReCC to rewrite the current query, as shown in Table 4. This allows us to assess its effectiveness in handling multi-document conversational QA tasks. It can be observed that compared to the T5-base rewriter, MaFeRw achieves better generation performance on this task. This validates that the rewriter trained using our method possesses generalization capability across different conversational tasks.

Conclusion

This paper introduces MaFeRw, a novel query rewriting method that enhances RAG performance by incorporating multi-aspect feedbacks. We find that dense reward feedback can lead to more stable and satisfactory generation results than single reward, and design four rewards from multiple aspects, such as gold document, retrieved documents and ground truth. The scores of these reward models are combined to optimize the rewriter using PPO algorithm. Experimental results on two conversational RAG datasets show MaFeRw outperforms baselines in generation metrics, and further analysis validates MaFeRw’s stability and transferability. Future work involves document re-ranking to better match retrieved documents with contextual constraints and prompt reconstruction to capture users’ complex intents in multi-turn dialogues.

Acknowledgments

This work was funded by the National Natural Science Foundation of China (NSFC) under Grants No. 62406013, the Beijing Advanced Innovation Center Funds for Future Blockchain and Privacy Computing(GJJ-23-006) and the Fundamental Research Funds for the Central Universities.

References

- Adlakha, V.; Dhuliawala, S.; Suleman, K.; de Vries, H.; and Reddy, S. 2022. TopiOCQA: Open-domain Conversational Question Answering with Topic Switching. *Transactions of the Association for Computational Linguistics*, 468–483.
- Anantha, R.; Vakulenko, S.; Tu, Z.; Longpre, S.; Pulman, S.; and Chappidi, S. 2020. Open-Domain Question Answering Goes Conversational via Question Rewriting. *arXiv: Information Retrieval, arXiv: Information Retrieval*.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In Goldstein, J.; Lavie, A.; Lin, C.-Y.; and Voss, C., eds., *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 65–72. Ann Arbor, Michigan: Association for Computational Linguistics.
- Bradley, R. A.; and Terry, M. E. 1952. Rank Analysis of Incomplete Block Designs: I. The Method of Paired Comparisons. *Biometrika*, 39(3/4).
- Chen, Z.; Zhao, J.; Fang, A.; Fetahu, B.; Rokhlenko, O.; and Malmasi, S. 2022. Reinforced Question Rewriting for Conversational Question Answering. In Li, Y.; and Lazaridou, A., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: Industry Track*, 357–370. Abu Dhabi, UAE: Association for Computational Linguistics.
- Choi, E.; He, H.; Iyyer, M.; Yatskar, M.; Yih, W.-t.; Choi, Y.; Liang, P.; and Zettlemoyer, L. 2018. QuAC: Question Answering in Context. In Riloff, E.; Chiang, D.; Hockenmaier, J.; and Tsujii, J., eds., *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2174–2184. Brussels, Belgium: Association for Computational Linguistics.
- Craswell, N. 2009. *Mean Reciprocal Rank*, 1703–1703. Boston, MA: Springer US. ISBN 978-0-387-39940-9.
- Dalton, J.; Xiong, C.; and Callan, J. 2020. TREC CAsT 2019: The Conversational Assistance Track Overview. arXiv:2003.13624.
- Du, X.; and Ji, H. 2022. Retrieval-Augmented Generative Question Answering for Event Argument Extraction. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 4649–4666. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Fang, Y.; Zhang, H.; Chen, H.; Ding, Z.; Long, B.; Lan, Y.; and Zhou, Y. 2022. From spoken dialogue to formal summary: An utterance rewriting for dialogue summarization. In *Proceedings of the 2022 conference of the North American chapter of the association for computational linguistics: human language technologies*, 3859–3869.
- Johnson, J.; Douze, M.; and Jegou, H. 2021. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 535–547.
- Kim, G.; Kim, H.; Park, J.; and Kang, J. 2021. Learn to Resolve Conversational Dependency: A Consistency Training Framework for Conversational Question Answering. In Zong, C.; Xia, F.; Li, W.; and Navigli, R., eds., *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 6130–6141. Online: Association for Computational Linguistics.
- Kwiatkowski, T.; Palomaki, J.; Redfield, O.; Collins, M.; Parikh, A.; Alberti, C.; Epstein, D.; Polosukhin, I.; Devlin, J.; Lee, K.; Toutanova, K.; Jones, L.; Kelcey, M.; Chang, M.-W.; Dai, A. M.; Uszkoreit, J.; Le, Q.; and Petrov, S. 2019. Natural Questions: A Benchmark for Question Answering Research. *Transactions of the Association for Computational Linguistics*, 7: 452–466.
- Lee, H.; Phatale, S.; Mansoor, H.; Mesnard, T.; Ferret, J.; Lu, K.; Bishop, C.; Hall, E.; Carbune, V.; Rastogi, A.; and Prakash, S. 2023. RLAI: Scaling Reinforcement Learning from Human Feedback with AI Feedback. arXiv:2309.00267.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; tau Yih, W.; Rocktäschel, T.; Riedel, S.; and Kiela, D. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. arXiv:2005.11401.
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, 74–81. Barcelona, Spain: Association for Computational Linguistics.
- Lin, S.; Yang, J.; Nogueira, R. F.; Tsai, M.; Wang, C.; and Lin, J. 2020. Conversational Question Reformulation via Sequence-to-Sequence Architectures and Pretrained Language Models. *CoRR*, abs/2004.01909.
- Ma, X.; Gong, Y.; He, P.; Zhao, H.; and Duan, N. 2023. Query Rewriting in Retrieval-Augmented Large Language Models. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 5303–5315. Singapore: Association for Computational Linguistics.
- Mao, K.; Dou, Z.; Liu, B.; Qian, H.; Mo, F.; Wu, X.; Cheng, X.; and Cao, Z. 2023a. Search-Oriented Conversational Query Editing. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023*, 4160–4172. Toronto, Canada: Association for Computational Linguistics.
- Mao, K.; Dou, Z.; Mo, F.; Hou, J.; Chen, H.; and Qian, H. 2023b. Large Language Models Know Your Contextual Search Intent: A Prompting Framework for Conversational Search. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*, 1211–1225. Singapore: Association for Computational Linguistics.

- Mo, F.; Mao, K.; Zhu, Y.; Wu, Y.; Huang, K.; and Nie, J.-Y. 2023. ConvGQR: Generative Query Reformulation for Conversational Search. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4998–5012. Toronto, Canada: Association for Computational Linguistics.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In Isabelle, P.; Charniak, E.; and Lin, D., eds., *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Qian, H.; and Dou, Z. 2022. Explicit Query Rewriting for Conversational Dense Retrieval. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 4725–4737. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- 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*. Association for Computational Linguistics.
- Schick, T.; Dwivedi-Yu, J.; Dessì, R.; Raileanu, R.; Lomeli, M.; Zettlemoyer, L.; Cancedda, N.; and Scialom, T. 2023. Toolformer: Language Models Can Teach Themselves to Use Tools. *ArXiv*, abs/2302.04761.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *arXiv: Learning*, *arXiv: Learning*.
- Sharma, S.; Yoon, D. S.; Dernoncourt, F.; Sultania, D.; Bagga, K.; Zhang, M.; Bui, T.; and Kotte, V. 2024. Retrieval Augmented Generation for Domain-specific Question Answering. *arXiv*:2404.14760.
- Siriwardhana, S.; Weerasekera, R.; Wen, E.; Kaluarachchi, T.; Rana, R.; and Nanayakkara, S. 2023. Improving the Domain Adaptation of Retrieval Augmented Generation (RAG) Models for Open Domain Question Answering. *Transactions of the Association for Computational Linguistics*, 11: 1–17.
- Su, H.; Jiang, S.; Lai, Y.; Wu, H.; Shi, B.; Liu, C.; Liu, Q.; and Yu, T. 2024. ARKS: Active Retrieval in Knowledge Soup for Code Generation. *arXiv*:2402.12317.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; Bikel, D.; Blecher, L.; Ferrer, C. C.; Chen, M.; Cucurull, G.; Esiobu, D.; Fernandes, J.; Fu, J.; Fu, W.; Fuller, B.; Gao, C.; Goswami, V.; Goyal, N.; Hartshorn, A.; Hosseini, S.; Hou, R.; Inan, H.; Kardas, M.; Kerkez, V.; Khabsa, M.; Kloumann, I.; Korenev, A.; Koura, P. S.; Lachaux, M.-A.; Lavril, T.; Lee, J.; Liskovich, D.; Lu, Y.; Mao, Y.; Martinet, X.; Mihaylov, T.; Mishra, P.; Molybog, I.; Nie, Y.; Poulton, A.; Reizenstein, J.; Rungta, R.; Saladi, K.; Schelten, A.; Silva, R.; Smith, E. M.; Subramanian, R.; Tan, X. E.; Tang, B.; Taylor, R.; Williams, A.; Kuan, J. X.; Xu, P.; Yan, Z.; Zarov, I.; Zhang, Y.; Fan, A.; Kambadur, M.; Narang, S.; Rodriguez, A.; Stojnic, R.; Edunov, S.; and Scialom, T. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *arXiv*:2307.09288.
- Wang, B.; Zheng, R.; Chen, L.; Liu, Y.; Dou, S.; Huang, C.; Shen, W.; Jin, S.; Zhou, E.; Shi, C.; Gao, S.; Xu, N.; Zhou, Y.; Fan, X.; Xi, Z.; Zhao, J.; Wang, X.; Ji, T.; Yan, H.; Shen, L.; Chen, Z.; Gui, T.; Zhang, Q.; Qiu, X.; Huang, X.; Wu, Z.; and Jiang, Y.-G. 2024. Secrets of RLHF in Large Language Models Part II: Reward Modeling. *arXiv*:2401.06080.
- Wang, Z.; Lin, P.; Zhang, H.; Chen, H.; Li, T.; Ding, Z.; Xu, S.; and Hu, J. 2023. Automatic Marketing Theme and Commodity Construction System for E-commerce. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Industry Track*, 501–508.
- Wu, Z.; Luan, Y.; Rashkin, H.; Reitter, D.; Hajishirzi, H.; Ostendorf, M.; and Tomar, G. S. 2022. CONQRR: Conversational Query Rewriting for Retrieval with Reinforcement Learning. In Goldberg, Y.; Kozareva, Z.; and Zhang, Y., eds., *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 10000–10014. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- Yang, H.; Yue, S.; and He, Y. 2023. Auto-GPT for Online Decision Making: Benchmarks and Additional Opinions. *arXiv*:2306.02224.
- Ye, F.; Fang, M.; Li, S.; and Yilmaz, E. 2023. Enhancing Conversational Search: Large Language Model-Aided Informative Query Rewriting. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*, 5985–6006. Singapore: Association for Computational Linguistics.
- Ye, L.; Lei, Z.; Yin, J.; Chen, Q.; Zhou, J.; and He, L. 2024. Boosting Conversational Question Answering with Fine-Grained Retrieval-Augmentation and Self-Check. *arXiv*:2403.18243.
- Zhang, H.; Lan, Y.; Pang, L.; Guo, J.; and Cheng, X. 2019. ReCoSa: Detecting the Relevant Contexts with Self-Attention for Multi-turn Dialogue Generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3721–3730.
- Zhang, J. 2023. Graph-ToolFormer: To Empower LLMs with Graph Reasoning Ability via Prompt Augmented by ChatGPT. *arXiv*:2304.11116.
- Zhao, P.; Zhang, H.; Yu, Q.; Wang, Z.; Geng, Y.; Fu, F.; Yang, L.; Zhang, W.; Jiang, J.; and Cui, B. 2024. Retrieval-Augmented Generation for AI-Generated Content: A Survey. *arXiv*:2402.19473.
- Ziegler, D. M.; Stiennon, N.; Wu, J.; Brown, T. B.; Radford, A.; Amodei, D.; Christiano, P.; and Irving, G. 2020. Fine-Tuning Language Models from Human Preferences. *arXiv*:1909.08593.