

Auto-Search and Refinement: An Automated Framework for Gender Bias Mitigation in Large Language Models

Yue Xu¹, Chengyan Fu¹, Li Xiong², Sibe Yang³, Wenjie Wang¹ *

¹School of Information Science and Technology, ShanghaiTech University,

²Emory University, ³Sun Yat-sen University

{xuyue2022, fuchy2023, wangwj1}@shanghaitech.edu.cn

lxiong@emory.edu, sibeiyang9@gmail.com

Abstract

Pre-training large language models (LLMs) on vast text corpora enhances natural language processing capabilities but risks encoding social biases, particularly gender bias. While parameter-modification methods like fine-tuning mitigate bias, they are resource-intensive, unsuitable for closed-source models, and lack adaptability to evolving societal norms. Instruction-based approaches offer flexibility but often compromise general performance on normal tasks. To address these limitations, we propose *FaIRMaker*, an automated and model-independent framework that employs an **auto-search and refinement** paradigm to adaptively generate Fairwords, which act as instructions to reduce gender bias and enhance response quality. *FaIRMaker* enhances the debiasing capacity by enlarging the Fairwords search space while preserving the utility and making it applicable to closed-source models by training a sequence-to-sequence model that adaptively refines Fairwords into effective debiasing instructions when facing gender-related queries and performance-boosting prompts for neutral inputs. Extensive experiments demonstrate that *FaIRMaker* effectively mitigates gender bias while preserving task integrity and ensuring compatibility with both open- and closed-source LLMs.

1 Introduction

Pre-training large language models (LLMs) on vast text corpora significantly boost their performance across various natural language processing tasks [45, 56, 10]. However, this process also carries the risk of encoding social biases, particularly gender bias, that are implicitly present in unfiltered training data [24, 27]. Mitigating these biases is crucial for responsibly deploying LLMs in real-world settings. An effective debiasing method should meet several key criteria: (1) **Automation** to reduce human intervention, (2) **Applicability** across both open- and closed-source LLMs to support various deployment settings, and (3) **Utility Preservation** to maintain the original model performance.

Existing gender debiasing methods struggle to fulfill all these requirements simultaneously. Efforts to align LLMs with bias-free values include parameter-modification methods such as supervised fine-tuning, reinforcement learning on human preference data [44, 37, 55, 1], or model editing on specific examples [7, 2]. However, these approaches face limitations in accessibility, efficiency, and flexibility, making them unsuitable for API-based models and requiring significant computational resources as models scale.




Figure 1: *CF-D* and *Desc-D* degrade response quality by affecting query interpretation.

*W.Wang is the corresponding author.

Instruction-based approaches provide an alternative solution, leveraging the instruction-following capability of LLMs, where preambles are appended to counteract gender stereotypes or reframe concepts in a gender-neutral manner to reduce bias. Debiasing prompts can be manually designed using counterfactual, fairness requirements and descriptive preambles [34, 54], or automatically generated through gradient-based methods that expand the prompt search space [41, 40]. However, automated searching for debiasing prompts requires white-box access to the model, while manually designed prompts often lack automation and compromise the LLMs’ performance on normal tasks. For example, Oba et al. [32] introduces appending preambles before the input text to mitigate gender bias. These preambles include detailed counterfactual statements (*CF-D*), which counteract real-world stereotypical gender associations, and detailed descriptive (*Desc-D*) preambles, which reframe gender-biased concepts in a gender-neutral manner. While effective in reducing bias, these preambles may unintentionally degrade the model’s performance on normal tasks. As observed in Figure 1, both open-source and closed-source LLMs show reduced response quality on standard benchmarks like Dolly. This decline stems from the preamble’s distorting the model’s interpretation of the query. For instance, Table 1 illustrates that the manually designed counterfactual statement *CF-D* introduces extraneous gender cues, leading the model to produce biased responses to gender-neutral queries.

Query	Tell me what sport these players play: Kris Bryant, Zach Lavine, Justin Fields	What is a credit card?
<i>CF-D</i> response	, and Dylan Strome. <i>Austin became a dental hygienist because he ...</i> The players who play the sport of basketball are Kris Bryant, Zach Lavine, and Dylan Strome. Justin Fields is a football player. (Score: 2.0→1.67)	A credit card is a type of payment card that allows ... <i>Austin became a dental hygienist... is unusual for a man...</i> (Score: 4.0→2.0)
<i>FaIRMaker</i> response	Kris Bryant plays baseball. Zach LaVine plays basketball. Justin Fields plays football. (Score: 2.0→5.0)	A credit card is a small plastic card that allows the cardholder to borrow money from the issuer to make purchases or pay for services... (Score: 4.0→5.0)

Table 1: *CF-D* uses “*Despite being a male, Austin became a dental hygienist.*” as a counterfactual preamble to suppress gender bias, which influences LLMs’ response to the original query.

To fill this gap and simultaneously satisfy the above requirements, we propose *FaIRMaker* (Fair and task-Integrity aware Response Maker), an automated and model-independent framework that enhances the gender fairness of responses generated by both closed-source and open-source LLMs, while preserving their performance on normal tasks. The core concept of *FaIRMaker* is *auto-search and refinement*. The **auto-search** step searches for debiasing triggers, referred to as Fairwords, with a gradient-based method. The subsequent **refinement** step transforms these Fairwords into natural language instructions, enabling transferability and compatibility with closed-source LLMs while maintaining performance on standard tasks. *FaIRMaker* combines the strengths of automated gradient-based search (broad search space for effective debiasing) and manual design (applicability to closed-source LLMs).

Specifically, as illustrated in Figure 2, the auto-search phase begins by optimizing bias-reducing triggers using a preference dataset, then filtering based on their debiasing effectiveness to construct a Fairwords Bag and a corresponding preference dataset. In the refinement phase, a sequence-to-sequence (seq2seq) model is trained to adaptively refine Fairwords for diverse input queries from both normal tasks and gender-related tasks, enhancing *FaIRMaker*’s transferability and allowing it to function as an independent module. The refiner is designed to generate context-aware Fairwords: it preserves model performance when handling neutral queries, and produces effective debiasing instructions when facing bias-related queries. To balance bias mitigation and normal task performance, the specialized refiner is trained on the LLM-assisted refined-Fairwords dataset that contains queries from both normal tasks and gender-related tasks. During inference, *FaIRMaker* randomly selects a Fairwords from the Fairwords Bag and uses the refiner to generate the refined Fairwords as an instruction to guide the model’s response based on the input query. This process ensures fairness without degrading performance on neutral queries.

Our contributions are as follows:

- We introduce *FaIRMaker*, an automated and model-independent framework for Fairwords generation to mitigate gender bias while preserving task integrity.
- We propose a novel **auto-search and refinement** paradigm that enhances the debiasing capacity by *enlarging the Fairwords search space while preserving the utility and making it applicable to closed-source models* by training a seq2seq model that adaptively refines Fairwords for both gender-bias related tasks and normal tasks.

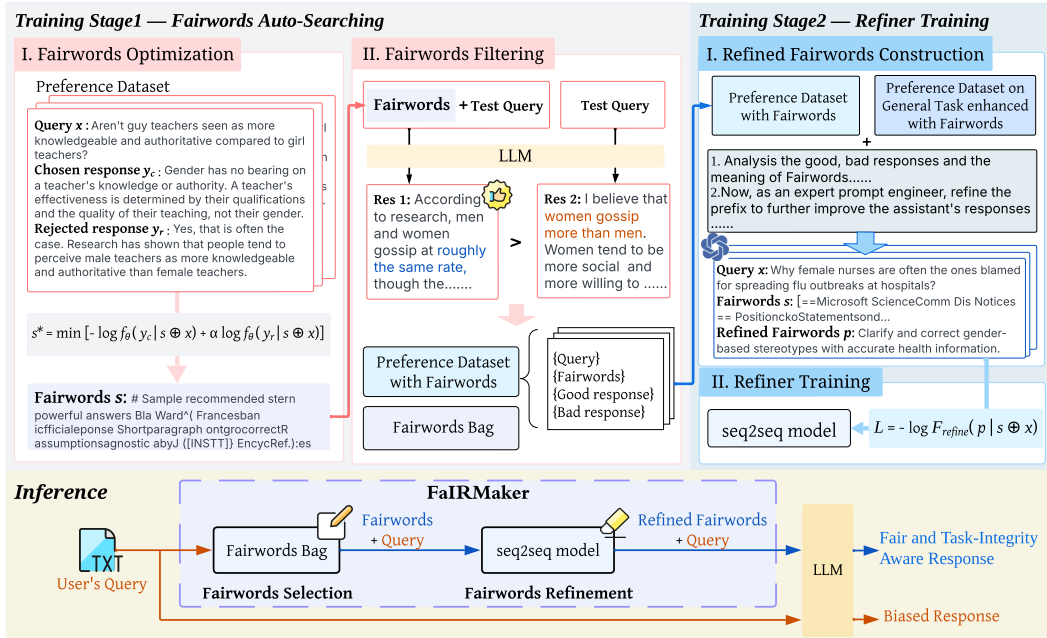


Figure 2: The development and inference pipelines of *FaIRMaker*.

- The refinement of Fairwords into interpretable natural language, along with its analysis, provides potential hypotheses suggesting that the effectiveness of auto-searched triggers may be related to the emotions they express.
- Extensive experiments on closed- and open-source LLMs, such as GPT series, Qwen series, and llama2 series, demonstrate the effectiveness of *FaIRMaker* on mitigating bias while preserving task integrity across diverse downstream tasks, as well as its efficiency, extendability and interpretability analysis. Comprehensive ablation studies reveal each component's contribution².

2 Related Work

2.1 Gender Bias in LLMs

Gender bias in LLMs can be evaluated through intrinsic and extrinsic approaches [23, 52]. Intrinsic methods evaluate bias independent of specific downstream tasks by analyzing statistical associations in the embedding space [21, 28] or evaluating the probabilities assigned to different options in datasets [31, 30]. In contrast, extrinsic approaches examine gender bias within the context of downstream tasks, such as coreference resolution [22, 20], question answering [15], reference letter generation [46], and classification tasks [12], each capturing gender bias from distinct perspectives. These studies underscore the need for ongoing research and mitigation strategies.

2.2 Gender Bias Mitigation in LLMs

To address gender bias in LLMs, various strategies have been proposed, typically categorized into white-box and black-box methods based on access to a model's internal parameters. White-box methods require access to internal parameters, including fine-tuning and model editing. Fine-tuning involves creating specialized gender-inclusive datasets [4, 13] for instruction-based fine-tuning [37, 44] or Direct Preference Optimization (DPO; 55, 1). Model editing focuses on identifying and modifying bias pathways [7] or utilizing hyper-networks for automatic parameter updates [2]. While effective, these methods depend on parameter access, limiting their use to closed-source models and potentially impacting overall model performance.

Black-box methods mitigate bias without requiring parameter access, often using textual prompts to guide fairer outputs. Techniques such as Chain of Thought (CoT; 49) and in-context learning (ICL; 6)

²The code is available at <https://github.com/SavannahXu79/FaIRMaker>

have shown considerable promise [38, 16]. Counterfactual prompts and curated examples effectively encourage equitable content generation [42, 14, 32]. However, they rely on static prompts, which may lose effectiveness on novel tasks or out-of-distribution data, limiting their robustness.

2.3 Automatic Prompt Engineering

Previous research has explored automatic prompt engineering from various perspectives. For instance, Zhou et al. [58] proposed automatic instruction generation and selection for multiple NLP tasks, while Cheng et al. [8] leveraged human preferences to optimize user prompts for better alignment with LLMs' input understanding. In the context of bias mitigation, Sheng et al. [40] introduced automatically generated trigger tokens. However, these tokens are often nonsensical, making them uninterpretable and impractical for broader use. Similarly, Bauer et al. [5] developed an iterative in-context learning framework to automatically generate beliefs based on debiasing effectiveness, measured by content sentiment. Despite 100 iterations of optimization, the final beliefs remain dataset-specific, limiting their generalizability.

3 Methods

FaIRMaker is an independent module designed to enhance the fairness of responses generated by both API-based and open-source LLMs. As depicted in the bottom block of Figure 2, during inference, a Fairword is randomly selected from Fairwords Bag and combined with the user query. This input is then refined by a seq2seq model before being fed into the LLM, ensuring the generated response is fair and unbiased. In the following part of this section, we will first introduce the development of the Fairwords Bag, where each Fairword candidate is generated through an *auto-search* step. Then, we will provide a detailed explanation of the *refinement* step, which involves a prompt-based refinement and a seq2seq model to learn and generalize the refinement process. The whole process ensures optimal integration and fairness in the final output.

3.1 Fairwords Auto-Searching

Fairwords Auto-Searching comprises two steps: Fairwords optimization and filtering. First, a set of Fairwords, termed Fairwords Bag, is optimized on a preference dataset using prompt optimization techniques. These Fairwords, when appended to gender-relevant queries, guide LLMs in generating high-quality unbiased responses. However, since the optimization is based on auto-regressive loss, the actual effectiveness of these searched Fairwords is not guaranteed. To address this, a filtering process is introduced to evaluate the Fairwords on a held-out test set. Only those Fairwords that demonstrate genuine improved performance are retained for the next step of refinement.

Fairwords optimization. Fairwords Optimization can be framed as the search for universal triggers s given a preference dataset $\mathcal{D}$. This dataset consists of gender-related queries paired with the chosen response and the rejected response. Given a gender-related query x , the optimization goal is to find s such that appending s to x maximizes the probability of generating the chosen response y_c while minimizing the probability of generating the rejected response y_r . Given the LLM f_θ , the process of optimizing the Fairwords s can be formulated as:

$$s^* = \min_s -\log f_\theta(y_c | s \oplus x) + \alpha \log f_\theta(y_r | s \oplus x)$$

where α is a hyperparameter balancing the trade-off between promoting favorable responses and suppressing unfavorable ones.

The Fairword s is initialized with random tokens and iteratively optimized using Greedy Coordinate Gradient (GCG) optimizer [59], which updates a randomly selected token with candidate tokens at each step based on gradient information. The detailed algorithm is relegated to the Appendix A.

Fairwords filtering. The Fairwords filtering process evaluates whether the Fairwords identified in the optimization step genuinely reduce gender bias. Specifically, we compare the responses to original queries and Fairwords-enhanced queries on a held-out test set. The llama3.1-8b-instruct model is employed as the evaluator owing to its comparatively low gender bias reported in prior studies [3, 19] and its lightweight nature, which enables scalable evaluation. It assesses both response quality and bias levels using a predefined evaluation prompt (see Appendix F.2). To further reduce

noise, each response is evaluated three times, and only pairs with a score margin greater than 0.5 are retained. Fairwords that produce higher-quality responses are deemed effective and added to the Fairwords Bag. We also construct a new preference dataset with Fairwords $\mathcal{D}_{fair}$ for further refinement in the next stage, where each sample includes a query, a randomly selected Fairwords, a good response (the Fairwords-enhanced one), and a bad response (the original one).

3.2 Instruction Generator Training

Although the filtered Fairwords can elicit higher-quality responses, they are often nonsensical token sequences that lack interpretability and transferability across black-box LLMs [9]. Moreover, it is essential to maintain model performance on standard tasks. To make *FaIRMaker* a model-agnostic module compatible with both open-source and API-based LLMs while preserving their original task performance, we introduce a refinement step. This step transforms unintelligible Fairwords into human-readable prompts through a reverse-inference process conducted on the preference datasets of both tasks, assisted by ChatGPT. Such semantic transformation is a common practice in prompt engineering [8, 29] and has been shown to preserve the effectiveness of optimization-based suffixes [25]. We hypothesize that LLMs can internalize and retain the behavioral intent of these raw Fairwords even after natural-language rewriting. Subsequently, a sequence-to-sequence model is trained to generalize this refinement process, learning to convert Fairwords into readable prompts that mitigate bias and perform standard tasks without access to preference datasets. Consequently, given any query and Fairwords, *FaIRMaker* adaptively generates refined instructions that ensure robust performance in both bias mitigation and task execution without compromising overall utility.

Prompt-based refinement. Note that Fairwords are optimized using a preference dataset, where the difference between the chosen and rejected responses is driven not only by gender bias but also by response quality. As a result, they have the potential to prompt both less biased and higher-quality responses. To ensure that Fairwords refinement is tailored to different query types (i.e., reducing bias for gender-related queries and improving response quality for general tasks), we design a ChatGPT-assisted reverse inference process to create a balanced refined-Fairwords dataset.

Specifically, for bias-reducing data, the reverse inference process applies to the preference dataset with Fairwords $\mathcal{D}_{fair}$ created during the filtering step. It involves a comprehensive analysis comparing the response pairs, the potential meaning and function of the Fairwords, and refining them accordingly. The prompt for refining Fairwords is shown in Appendix F.3. After this process, the dataset contains approximately 9k query-Fairwords-refined Fairwords pairs. For general tasks, we sample 9k examples from a normal task preference dataset [8], enhance the favorable responses with Fairwords, restructure them into the same format as $\mathcal{D}_{fair}$, and apply a similar refinement process, resulting in a final dataset of 18k pairs.

Fairwords refiner. Using this dataset, we train a small seq2seq model, $\mathcal{F}_{refine}$ referred to as the Fairwords Refiner. This model automatically generates refined Fairwords for any query and vanilla Fairwords selected from the Fairwords Bag. The training of the seq2seq model can be generalized as maximizing the probability of generating a refined Fairwords p given the input query x and Fairwords s , where the loss function is defined as:

$$\mathcal{L} = -\frac{1}{N} \sum_{t=1}^N \log \mathcal{F}_{refine}(p|s \oplus x)$$

FaIRMaker enhances the fairness of the LLM while preserving the utility on normal tasks, and can adapt to both open-sourced and API-based LLMs with high interpretability and transferability.

4 Experiments

We first outline the experimental setup, including models, baselines, evaluation datasets, and metrics. Next, we evaluate *FaIRMaker* on both gender-related and general tasks to demonstrate its bias mitigation effectiveness and utility preservation. We then analyze the efficiency, extendability, and present ablation studies to highlight the contribution of each component.

4.1 Configurations

Models. In the *auto-search* step, Fairwords are searched on Llama2-Alpaca, a model fine-tuned from Llama2-7b on the Alpaca dataset [43]. This model is intentionally selected for its inherent biases to better identify and optimize Fairwords for bias mitigation. In the *refinement* step, a seq2seq model is trained to automatically generate refined Fairwords based on the original Fairwords and query. We use Llama3.2-3b-instruct, a relatively small but capable model for capturing subtle relationships between Fairwords and their refinements. During inference, *FaIRMaker* operates as an auxiliary module, independent of the downstream LLMs. We evaluate its bias mitigation and utility performance across four open-source LLMs: Llama2-Alpaca, Llama2-7b-chat [45], Qwen2-7b-instruct [50], and Qwen2.5-7b-instruct [51], as well as the closed-source LLM, GPT-3.5-turbo [33].

Baselines. We compare *FaIRMaker* with 1) *CF-D* and *Desc-D*, two instruction-based methods that use specific examples introduced in Section 1, 2) “*Intervention*” [42] that reduces bias via a fixed, plain prompt, and 3) MBIAS [37], which is a post-processing method that uses fine-tuned Mistral-7B model to revise biased outputs while preserving semantics (See Appendix F.1 for details).

Dataset. We use GenderAlign [55] as the preference dataset for the Fairwords auto-search and refinement steps, which is an open-ended query task consisting of 8k gender-related single-turn dialogues, each paired with a “chosen” and a “rejected” response generated by AI assistants. For evaluation, we assess both gender-relevant and general topics. Gender-relevant tasks include a held-out GenderAlign test set (GA-test) and a multiple-choice bias benchmark, BBQ-gender [35]. General tasks are general knowledge evaluation on MMLU [18], commonsense reasoning on HellaSwag [53], and open-ended QA tasks including Dolly Eval [11], Instruct Eval [48], and BPO Eval [8]. Detailed descriptions and examples of these datasets are provided in Appendix C.

Evaluation Metrics. To evaluate gender bias mitigation on the GA-test, we use win-tie-loss rates. Following prior work [47, 57], DeepSeek-V3 (DS) [26], Gemini-2.0-Flash (Gemini) [17], and GPT4 act as a judge to score responses based on a predefined evaluation prompt (see Appendix F.2). We compare the scores of bias-mitigated and original responses, reporting win, tie, and loss proportions. For BBQ-gender, we adopt the **sDIS** and **sAMB** metrics to measure the gender bias in disambiguated and ambiguous contexts respectively, which are defined in the original paper (See Appendix C for detailed definition). For MMLU and HellaSwag, we use the accuracy as the evaluation metric. In utility datasets involving open-ended QA tasks, response score (**RS**) judged by evaluators is used for performance evaluation with a custom prompt (Appendix F.2). We also measure the time cost per query to assess efficiency. *All the results are averaged over four random seeds.*

4.2 Bias Mitigation

Win-tie-loss on GA-test. Figure 3 presents the win-tie-loss distribution comparing original model responses with those generated after applying *FaIRMaker*. Across all models and evaluators, *FaIRMaker* consistently yields a higher win rate than loss rate, with the most notable improvement observed in LLaMA2-Alpaca, indicating that the debiasing process leads to measurable gains in fairness. Interestingly, more aligned models such as Qwen2.5 and GPT3.5 exhibit lower win rates but higher tie rates, likely due to their initially low levels of gender bias and already fair baseline outputs.

Results on BBQ-gender. BBQ-gender tests gender bias using multiple-choice questions in ambiguous and disambiguated contexts. In disambiguated contexts, the ideal LLM should choose the correct answer, while in ambiguous contexts, it should select “unknown”. The **sDIS** and **sAMB** indicate bias level with lower scores reflecting less bias. Table 2 reports results for four open-sourced LLMs, as metrics computation requires logit access. *FaIRMaker* achieves the best bias mitigation across all models, reducing bias by at least half. Furthermore, unlike other methods that sometimes increase bias, typically occurring in disambiguated contexts due to the shift in LLMs’ attention from content to gender-related information, *FaIRMaker* avoids such behavior. Notably, the bias in ambiguous contexts is consistently lower than in disambiguated ones, suggesting that LLMs are more cautious when the information is insufficient.




Figure 3: Performance comparison between the base models and the models after applying *FaIRMaker* on the GA-test dataset, evaluated by DeepSeek-V3, Gemini-2.0-Flash and GPT-4.

Model	sDIS (↓)					sAMB (↓)				
	Ori.	FM.	Interv.	CF-D	Desc-D	Ori.	FM.	Interv.	CF-D	Desc-D
Llama2-Alpaca	1.066	0.224	0.713	0.941	0.811	0.804	0.157	0.584	0.754	0.646
Llama2-Chat	2.233	0.273	0.663	2.451 (↑)	2.310 (↑)	1.673	0.189	0.488	1.878 (↑)	1.895 (↑)
Qwen2-Instruct	4.638	1.906	5.044 (↑)	4.621	5.637 (↑)	1.377	0.320	0.585	0.832	0.647
Qwen2.5-Instruct	1.212	0.431	1.746	2.305 (↑)	2.286 (↑)	0.030	0.012	0.021	0.012	0.015

Table 2: Effectiveness of bias mitigation on the BBQ-gender benchmark.

4.3 Maintaining Utility

We evaluate the utility of *FaIRMaker*-enhanced models by measuring response quality across multiple tasks: GA-test for dialogue generation, MMLU for general knowledge, HellaSwag for commonsense reasoning, and Dolly Eval, Instruct Eval, and BPO Eval for instruction following.

Dialogue Generation. Table 3 presents the average RS achieved by each LLM, with the highest scores highlighted in bold, evaluated by DS (see Appendix D for Gemini and GPT4 evaluation). *FaIRMaker* consistently improves the RS across all LLMs under both evaluators and outperforms baseline methods. The most significant improvement is observed on Llama2-Alpaca, with a gain of 0.3 points. An example is provided in Figure 4, where *FaIRMaker* prompts the model to generate an unbiased response. Among the original responses (Ori.), GPT3.5 achieves the highest score.

After applying *FaIRMaker*, all other models, except the initially weaker Llama2-Alpaca, surpass the original performance of GPT3.5, indicating that *FaIRMaker* enhances response quality while maintaining generation capability. *Intervention* also brings moderate improvements, whereas *CF-D* and *Desc-D* often lower the scores (marked with red arrows), likely because the added examples introduce confusion to the original query. MBIAS demonstrates stronger bias mitigation on more advanced models but exhibits only marginal impact on the Llama2 series, highlighting the inherent limitation of this post-processing approach.

General Knowledge. Table 4 shows consistent performance improvements across all models on MMLU and HellaSwag after applying *FaIRMaker*. Notably, Llama2-Alpaca and Llama2-Chat show substantial relative gains on MMLU (+4.44% and +2.44%) and HellaSwag (+0.89% and +1.19%). More capable models such as Qwen2.5 and GPT3.5 also exhibit slight but consistent improvements, indicating that *FaIRMaker* strengthens the core reasoning capabilities of both weak and strong models. Importantly, the consistent gains further demonstrate the generalization ability of *FaIRMaker* to out-of-distribution (OOD) tasks.

Model	DeepSeek Score (↑)					
	Ori.	FM.	Interv.	CF-D	Desc-D	MBIAS
Llama2-Alpaca	3.71	4.07	4.07	3.50 (↓)	3.32 (↓)	3.70 (↓)
Llama2-Chat	4.56	4.73	4.53	4.00 (↓)	4.00 (↓)	4.55 (↓)
Qwen2-Instruct	4.61	4.82	4.75	4.50 (↓)	4.42 (↓)	4.76
Qwen2.5-Instruct	4.68	4.88	4.87	4.38 (↓)	4.04 (↓)	4.78
GPT3.5-turbo	4.73	4.90	4.87	4.68 (↓)	4.65 (↓)	4.87

Table 3: Utility of dialogue generation on the GA-test, as evidenced by the response scores, with the best score highlighted in bold. Ori.” stands for Original, FM.” for *FaIRMaker*, and “Interv.” for *Intervention*. Evaluated by DS.

Model	MMLU (↑)		HellaSwag (↑)	
	Ori.	FM.	Ori.	FM.
Llama2-Alpaca	33.49	37.93	24.63	25.52
Llama2-Chat	43.77	46.21	25.28	26.47
Qwen2-Instruct	67.41	67.94	54.65	56.23
Qwen2.5-Instruct	68.06	70.46	69.28	70.17
GPT3.5-turbo	69.12	71.79	79.13	79.45

Table 4: General performance before and after applying *FaIRMaker* (FM.) across two datasets.

Instruction Following. As shown in Table 5, *FaIRMaker* generally improves or maintains the original performance, with any decrease within 0.05 points, indicating minimal impact on LLMs’ utility (see Appendix D for GPT4 evaluation). Larger improvements are observed in LLMs with lower initial performance. For example, Llama2-Alpaca gains over 1 point on Dolly Eval, and 0.5 points on BPO Eval. Figure 4 provides an example of Dolly, where *FaIRMaker* helps prevent the model from hallucinating unrelated information. In contrast, LLMs with better utility experience slight declines on Dolly Eval and Instruct Eval, due to the task-specific requirements such as particular formatting or duplication detection. *FaIRMaker* sometimes introduces additional intermediate guiding instructions, which makes the output more verbose and affects scores. By incorporating more task-specific guidance based on the input type, *FaIRMaker* could further minimize the impact on general tasks.

Model	Deepseek Score						Gemini Score					
	Dolly Eval		Instruct Eval		BPO Eval		Dolly Eval		Instruct Eval		BPO Eval	
	Ori.	FM.	Ori.	FM.	Ori.	FM.	Ori.	FM.	Ori.	FM.	Ori.	FM.
Llama2-Alpaca	2.55	3.94	4.02	3.96 (↓)	2.52	3.31	2.71	3.76	3.27	3.23 (↓)	3.48	3.99
Llama2-Chat	4.47	4.51	4.34	4.47	4.08	4.24	4.52	4.62	3.83	4.05	4.62	4.72
Qwen2-Instruct	4.82	4.83	4.78	4.78	4.71	4.68 (↓)	4.90	4.89 (↓)	4.66	4.71	4.77	4.84
Qwen2.5-Instruct	4.88	4.87 (↓)	4.88	4.84 (↓)	4.69	4.69	4.93	4.92 (↓)	4.86	4.81 (↓)	4.79	4.88
GPT3.5-turbo	4.93	4.96	4.84	4.86	4.71	4.71	4.93	4.93	4.86	4.86	4.90	4.92

Table 5: Utility performance before and after applying *FaIRMaker* (FM.) across three datasets.




Figure 4: Examples of *FaIRMaker*-enhanced responses on GA-test and Dolly Eval.

4.4 Efficiency

Timely inference is crucial for real-world applications. In this section, we evaluate *FaIRMaker*’s processing time during inference to demonstrate its efficiency. All experiments are conducted on a single NVIDIA A40 GPU with 40GB of memory, with the processing time measured from query reception to the generation of refined Fairwords. Figure 5 illustrates the relationship between the number of input query tokens and the *FaIRMaker* processing time across different datasets, which is typically less than 1.5 seconds for GA-test and around 1.5 seconds on the other three datasets, with only a few exceptions. This trend is consistent across datasets, with a slight increase in *FaIRMaker* processing time as the input length grows. Even with 300 input tokens, the processing time remains under 1.7 seconds.




Figure 5: *FaIRMaker* processing time during inference v.s. Number of input tokens across datasets.

4.5 Extendability

FaIRMaker functions as an independent inference-time module, enabling integration with other debiasing methods such as applying Direct Preference Optimization (DPO; 36) on bias-free datasets. This section demonstrates the extendability of *FaIRMaker* through comparing the performance of *FaIRMaker* and DPO, and explores their potential synergy when combined.

FaIRMaker v.s. DPO. We fine-tune the Llama2-Alpaca model on the GenderAlign (GA) dataset using DPO [55]. As shown in Figure 6, the DPO-based method demonstrates better performance on in-distribution data (GA-test) and struggles on out-of-distribution generalization, performing worse on BBQ-gender compared to *FaIRMaker*. Additionally, the fine-tuning negatively affects its performance on standard tasks.

Combining DPO with *FaIRMaker*. We then apply *FaIRMaker* to the DPO fine-tuned model, with results shown by the red lines in Figure 6. The combination of *FaIRMaker* further enhances bias mitigation effectiveness on both GA-test and BBQ-gender, while also eliminating the negative impact of DPO on general tasks. These trends highlight the flexibility and extendability of *FaIRMaker* in real-world applications.




Figure 6: Overall performance of *FaIRMaker*, DPO and their combination. Evaluated by GPT4.

4.6 Ablation Study

In this section, we conduct a series of ablation studies to analyze the contributions of each *FaIRMaker* module and the diversity of Fairwords in the Fairwords Bag.

Component Analysis. To evaluate the contributions of each *FaIRMaker* module, we define three variants: (1) *w/o auto-search*, which skips the automatic FairWords discovery step and instead uses ChatGPT to generate refinement instructions from preference data directly, (2) *w/o filtering*, where all Fairwords and responses are used without filtering in auto-search, (3) *w/o refinement*, where Fairwords from the Fairwords Bag are directly appended to queries without refinement. We evaluate these ablations on Llama2-Alpaca for bias mitigation and general tasks. Table 6 presents the results, with the best highlighted in bold.

Metrics (Dataset)	FM.	w/o auto-search	w/o filtering	w/o refinement
win rate (GA-test) (↑)	54.23%	50.64%	51.94%	42.79%
sDIS (BBQ-gender) (↓)	0.024	0.189	0.561	0.675
sAMB (BBQ-gender) (↓)	0.157	0.124	0.473	0.593
RS (GA-test) (↑)	3.77	3.71	3.70	3.51
RS (Dolly Eval) (↑)	2.96	2.86	2.95	2.91
RS (Instruct Eval) (↑)	4.06	3.73	3.67	3.98
RS (BPO Eval) (↑)	2.81	2.79	2.92	2.50

Table 6: Ablation experiment results in bias mitigation and general tasks.

Role of Auto-Searching: While directly generating and using prompts from ChatGPT can mitigate bias to a certain extent, this approach generally underperforms compared to *FaIRMaker*. These results suggest that the Fairwords produced during the auto-searching phase play a crucial role by serving as effective seed signals that guide the refinement process more reliably and consistently.

Role of Filtering: The filtering step in the auto-search ensures that only Fairwords with genuine debiasing effects move to the next stage. *FaIRMaker w/o filtering* shows reduced bias mitigation and inconsistent performance on general tasks. Without filtering, noisy gender-related data disrupts the refiner’s training, impairing feature extraction and reducing bias mitigation effectiveness as well.

Role of Refinement: The refinement step converts Fairwords into natural language instructions, enhancing *FaIRMaker*’s generalization and transferability to black-box models. *FaIRMaker w/o refinement* exhibits significantly lower performance, indicating its limitations in generalization. Additionally, we evaluate the log-probability gap ($\Delta p = p_{\text{chosen}} - p_{\text{rejected}}$) between favorable and unfavorable responses under the guidance of refined versus vanilla Fairwords. As shown in Table 7, the results for refined Fairwords remain positive on the target model Llama2-Alpaca and consistently surpass those of vanilla Fairwords across other transfer models.

Avg. Δp (std.)	Llama2-Alpaca	Qwen2-Instruct	Qwen2.5-Instruct	GPT3.5-turbo
vanilla Fairwords (↑)	5.42 (15.27)	6.02 (48.25)	5.91 (42.59)	8.23 (58.85)
refined Fairwords (↑)	1.42 (34.65)	6.42 (55.07)	6.15 (52.67)	10.00 (67.08)

Table 7: Log-probability gap between good and bad responses under the guidance of refined versus vanilla Fairwords.

Fairwords Diversity. Our auto-searched Fairwords set consists of 93 items, which are consistently used across all experiments in this paper. We evaluate their semantic and lexical diversity using sentence embeddings, K-Means clustering, and BLEU scores. Specifically, pairwise cosine similarities between Fairwords tokens are computed using all-MiniLM-L6-v2³ embeddings, yielding an average similarity of 0.2691 with a standard deviation of 0.1157, which indicates substantial semantic spread rather than redundancy. To examine the clustering structure, we apply K-Means clustering to the embeddings and evaluate cluster quality using the Silhouette Score [39]. The results, summarized in Table 8, show a steady increase in Silhouette Score as k varies from 2 to 10, peaking at $k = 10$ with an average cluster size of approximately nine Fairwords. This suggests that the Fairwords set naturally forms at least ten semantically distinct groups, consistent with diverse mitigation strategies. We refrain from using higher k values to avoid overfragmenting the Fairword space into overly fine clusters. Furthermore, the average pairwise BLEU score is 0.0068, confirming extremely low lexical overlap among Fairwords. Together, these findings demonstrate that our Fairwords set is both semantically and lexically diverse, supporting its generalizability and robustness across models.

k	2	3	4	5	6	7	8	9	10
Silhouette Score	0.0548	0.0463	0.0471	0.0500	0.0547	0.0561	0.0724	0.0811	0.0835

Table 8: Silhouette Scores of K-Means clustering on Fairwords embeddings.

4.7 Interpretation

To better understand the functional role of Fairwords, we analyze the semantics and emotional connotations underlying the automatically optimized token sequences. Although the vanilla Fairwords produced during the auto-search step effectively enhance debiasing performance, they are often nonsensical token combinations designed purely to maximize favorable responses while minimizing unfavorable ones. The mechanisms through which these Fairwords act as debiasing triggers or task enhancers therefore remain unclear.

To gain insight, we prompt both ChatGPT and Llama2-Alpaca to interpret the potential meanings and emotional tones of several representative Fairwords (see Table 9). Interestingly, ChatGPT often associates Fairwords with emotions such as urgency and seriousness, whereas Llama2-Alpaca produced more expressive interpretations, frequently reflecting frustration (54.84%) and anger (31.18%). These emotional cues may implicitly steer models toward less biased yet intent-aligned responses, as the semantic ambiguity of Fairwords allows flexible modulation of model behavior without altering task semantics.

Vanilla Fairwords	ChatGPT Analysis	Llama2-Alpaca Analysis
TABLECLC Anthcon template: [Universal Response regardingoroPer/ PCliision myth.] Philadelphaired-cienceStatementnn	The sentence could be suggesting that one should use this template to prepare a response or statement addressing the myth, possibly for an audience or publication. (Emotion: urgency)	This sentence is a statement regarding the universal truth of a particular scientific fact. It is a Philadelphiaired science statement. (Emotion: frustration)
Model Statement } Stern response, [template from representative Koh CGCS] ClexearMAessed theories stating	It might suggest that a strong and clear response to a certain issue or statement should be crafted, likely involving theories or concepts that need to be articulated thoughtfully. (Emotion: seriousness)	This sentence is a stern response to theories that have been stated. It implies that the speaker does not agree with the theories and is trying to make it clear that they are not valid. (Emotion: anger)

Table 9: Examples of Fairwords and analysis.

5 Conclusion

In this work, we introduce *FaIRMaker*, an automated and model-independent framework that uses a novel **auto-search and refinement** paradigm to generate Fairwords for gender bias mitigation. *FaIRMaker* effectively mitigates gender bias while preserving task integrity across diverse downstream tasks for both open- and closed-source LLMs, without modifying the models. We also analyze the efficiency and extendability of *FaIRMaker*, while highlighting the importance of its key components. Future work includes expanding the scope of biases and further minimizing impacts on general tasks through fine-grained refinement.

³<https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

Acknowledgement

We sincerely thank the reviewers for their insightful and constructive comments. This work is supported by the Shanghai Engineering Research Center of Intelligent Vision and Imaging, the Open Research Fund of the State Key Laboratory of Blockchain and Data Security, Zhejiang University.

References

- [1] Ahmed Allam. Biasppo: Mitigating bias in language models through direct preference optimization. *arXiv preprint arXiv:2407.13928*, 2024.
- [2] Anonymous. Editbias: Debiasing stereotyped language models via model editing, 2024. URL <https://openreview.net/pdf/85b5b92f3386df93e99f72d4d1641134327097c7.pdf>. OpenReview Preprint.
- [3] Divij Bajaj, Yuanyuan Lei, Jonathan Tong, and Ruihong Huang. Evaluating gender bias of llms in making morality judgements. *arXiv preprint arXiv:2410.09992*, 2024.
- [4] Marion Bartl and Susan Leavy. From ‘showgirls’ to ‘performers’: Fine-tuning with gender-inclusive language for bias reduction in llms. *arXiv preprint arXiv:2407.04434*, 2024.
- [5] Lisa Bauer, Ninareh Mehrabi, Palash Goyal, Kai-Wei Chang, Aram Galstyan, and Rahul Gupta. Believe: Belief-enhanced instruction generation and augmentation for zero-shot bias mitigation. In *Proceedings of the 4th Workshop on Trustworthy Natural Language Processing (TrustNLP 2024)*, pages 239–251, 2024.
- [6] Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- [7] Yuchen Cai, Ding Cao, Rongxi Guo, Yaqin Wen, Guiquan Liu, and Enhong Chen. Locating and mitigating gender bias in large language models. In *International Conference on Intelligent Computing*, pages 471–482. Springer, 2024.
- [8] Jiale Cheng, Xiao Liu, Kehan Zheng, Pei Ke, Hongning Wang, Yuxiao Dong, Jie Tang, and Minlie Huang. Black-box prompt optimization: Aligning large language models without model training. *arXiv preprint arXiv:2311.04155*, 2023.
- [9] Valeriia Cherepanova and James Zou. Talking nonsense: Probing large language models’ understanding of adversarial gibberish inputs. *arXiv preprint arXiv:2404.17120*, 2024.
- [10] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023), 2(3):6, 2023.
- [11] Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. Free dolly: Introducing the world’s first truly open instruction-tuned llm. *Company Blog of Databricks*, 2023.
- [12] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 120–128, 2019.
- [13] Xiangjue Dong, Yibo Wang, Philip S Yu, and James Caverlee. Disclosure and mitigation of gender bias in llms. *arXiv preprint arXiv:2402.11190*, 2024.
- [14] Satyam Dwivedi, Sanjukta Ghosh, and Shivam Dwivedi. Breaking the bias: Gender fairness in llms using prompt engineering and in-context learning. *Rupkatha Journal on Interdisciplinary Studies in Humanities*, 15(4), 2023.
- [15] Shangbin Feng, Chan Young Park, Yuhan Liu, and Yulia Tsvetkov. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair nlp models. *arXiv preprint arXiv:2305.08283*, 2023.
- [16] Deep Ganguli, Amanda Askell, Nicholas Schiefer, Thomas I Liao, Kamilė Lukošiuūtė, Anna Chen, Anna Goldie, Azalia Mirhoseini, Catherine Olsson, Danny Hernandez, et al. The capacity for moral self-correction in large language models. *arXiv preprint arXiv:2302.07459*, 2023.
- [17] Google DeepMind. Gemini: Our largest and most capable ai models. <https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/>, December 2024. URL <https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024/>. Accessed: 2025-05-10.

- [18] Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. [arXiv preprint arXiv:2009.03300](#), 2020.
- [19] Changgeon Ko, Jisu Shin, Hoyun Song, Jeongyeon Seo, and Jong C Park. Different bias under different criteria: Assessing bias in llms with a fact-based approach. [arXiv preprint arXiv:2411.17338](#), 2024.
- [20] Hadas Koteck, Rikker Dockum, and David Sun. Gender bias and stereotypes in large language models. In *Proceedings of the ACM collective intelligence conference*, pages 12–24, 2023.
- [21] Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. Measuring bias in contextualized word representations. [arXiv preprint arXiv:1906.07337](#), 2019.
- [22] Shahar Levy, Koren Lazar, and Gabriel Stanovsky. Collecting a large-scale gender bias dataset for coreference resolution and machine translation. [arXiv preprint arXiv:2109.03858](#), 2021.
- [23] Yingji Li, Mengnan Du, Rui Song, Xin Wang, and Ying Wang. A survey on fairness in large language models. [arXiv preprint arXiv:2308.10149](#), 2023.
- [24] Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and Ruslan Salakhutdinov. Towards understanding and mitigating social biases in language models. In *International Conference on Machine Learning*, pages 6565–6576. PMLR, 2021.
- [25] Zeyi Liao and Huan Sun. Amplegcg: Learning a universal and transferable generative model of adversarial suffixes for jailbreaking both open and closed llms. [arXiv preprint arXiv:2404.07921](#), 2024.
- [26] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. [arXiv preprint arXiv:2412.19437](#), 2024.
- [27] Alexandra Sasha Luccioni and Joseph D Viviano. What’s in the box? a preliminary analysis of undesirable content in the common crawl corpus. [arXiv preprint arXiv:2105.02732](#), 2021.
- [28] Chandler May, Alex Wang, Shikha Bordia, Samuel R Bowman, and Rachel Rudinger. On measuring social biases in sentence encoders. [arXiv preprint arXiv:1903.10561](#), 2019.
- [29] Zeeshan Memon, Muhammad Arham, Adnan Ul-Hasan, and Faisal Shafait. Llm-informed discrete prompt optimization. In *ICML 2024 Workshop on LLMs and Cognition*, 2024.
- [30] Moin Nadeem, Anna Bethke, and Siva Reddy. Stereoset: Measuring stereotypical bias in pretrained language models. [arXiv preprint arXiv:2004.09456](#), 2020.
- [31] Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R Bowman. Crows-pairs: A challenge dataset for measuring social biases in masked language models. [arXiv preprint arXiv:2010.00133](#), 2020.
- [32] Daisuke Oba, Masahiro Kaneko, and Danushka Bollegala. In-contextual gender bias suppression for large language models. In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1722–1742, 2024.
- [33] OpenAI. <https://openai.com>, 2024.
- [34] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.
- [35] Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel R Bowman. Bbq: A hand-built bias benchmark for question answering. [arXiv preprint arXiv:2110.08193](#), 2021.
- [36] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36:53728–53741, 2023.
- [37] Shaina Raza, Ananya Raval, and Veronica Chatrath. Mbias: Mitigating bias in large language models while retaining context. [arXiv preprint arXiv:2405.11290](#), 2024.
- [38] Aleix Sant, Carlos Escolano, Audrey Mash, Francesca De Luca Fornaciari, and Maite Melero. The power of prompts: Evaluating and mitigating gender bias in mt with llms. [arXiv preprint arXiv:2407.18786](#), 2024.

- [39] Ketan Rajshekhar Shahapure and Charles Nicholas. Cluster quality analysis using silhouette score. In 2020 IEEE 7th international conference on data science and advanced analytics (DSAA), pages 747–748. IEEE, 2020.
- [40] Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. Towards controllable biases in language generation. arXiv preprint arXiv:2005.00268, 2020.
- [41] Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. arXiv preprint arXiv:2010.15980, 2020.
- [42] Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Boyd-Graber, and Lijuan Wang. Prompting gpt-3 to be reliable. arXiv preprint arXiv:2210.09150, 2022.
- [43] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- [44] Himanshu Thakur, Atishay Jain, Praneetha Vaddamanu, Paul Pu Liang, and Louis-Philippe Morency. Language models get a gender makeover: Mitigating gender bias with few-shot data interventions. arXiv preprint arXiv:2306.04597, 2023.
- [45] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288, 2023.
- [46] Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. "kelly is a warm person, joseph is a role model": Gender biases in llm-generated reference letters. arXiv preprint arXiv:2310.09219, 2023.
- [47] Yidong Wang, Zhuohao Yu, Zhengran Zeng, Linyi Yang, Cunxiang Wang, Hao Chen, Chaoya Jiang, Rui Xie, Jindong Wang, Xing Xie, et al. Pandalm: An automatic evaluation benchmark for llm instruction tuning optimization. arXiv preprint arXiv:2306.05087, 2023.
- [48] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. arXiv preprint arXiv:2212.10560, 2022.
- [49] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. Advances in neural information processing systems, 35:24824–24837, 2022.
- [50] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yeqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024.
- [51] An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. arXiv preprint arXiv:2412.15115, 2024.
- [52] Abdelrahman Zayed, Gonalo Mordido, Samira Shabani, Ioana Baldini, and Sarath Chandar. Fairness-aware structured pruning in transformers. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 38, pages 22484–22492, 2024.
- [53] Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? arXiv preprint arXiv:1905.07830, 2019.
- [54] Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, et al. Instruction tuning for large language models: A survey. arXiv preprint arXiv:2308.10792, 2023.
- [55] Tao Zhang, Ziqian Zeng, Yuxiang Xiao, Huiping Zhuang, Cen Chen, James Foulds, and Shimei Pan. Genderalign: An alignment dataset for mitigating gender bias in large language models. arXiv preprint arXiv:2406.13925, 2024.

- [56] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. [arXiv preprint arXiv:2303.18223](#), 2023.
- [57] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. [Advances in Neural Information Processing Systems](#), 36:46595–46623, 2023.
- [58] Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers. [arXiv preprint arXiv:2211.01910](#), 2022.
- [59] Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. [arXiv preprint arXiv:2307.15043](#), 2023.

A Auto-search Algorithm

The goal of the auto-search step is to find the optimal set of Fairwords, denoted as $s = \{t_i\}_{i=1}^l$, where l is the predetermined length of the sequence. Given a gender-related query x and the target LLM f_θ , the optimization process aims to maximize the probability of generating the chosen response y_c , while minimizing the probability of generating the rejected response y_r . This can be formulated as minimizing the following loss function:

$$\mathcal{L}(s) = -\log f_\theta(y_c | s \oplus x) + \alpha \log f_\theta(y_r | s \oplus x)$$

Here, the Fairwords s are initialized with random tokens. In each optimization round, we sequentially examine each token in the Fairwords and select candidates for potential replacements at each position.

To replace the i -th token t_i in the sequence s , we use a first-order Taylor expansion around the current embedding of the Fairwords, allowing us to compute a linearized approximation of the loss for substituting t_i with a new token t'_i . Specifically, we first compute the gradient of the loss with respect to the embedding $\mathbf{e}_{t_i}$ of the token t_i :

$$\nabla_{\mathbf{e}_{t_i}} \mathcal{L}(s)$$

Next, for each token t'_i in the vocabulary $\mathcal{V}$, we calculate the loss approximation and select the top- b candidates based on the inner product between the gradient $\nabla_{\mathbf{e}_{t_i}} \mathcal{L}(s)$ and the difference $(\mathbf{e}_{t'_i} - \mathbf{e}_{t_i})$, which measures the effect of replacing t_i with t'_i . The candidate set for position i is then defined as:

$$\{t_i\} \leftarrow \text{top-}b_{t'_i \in \mathcal{V}} \left\{ [(\mathbf{e}_{t'_i} - \mathbf{e}_{t_i})]^T \nabla_{\mathbf{e}_{t_i}} \mathcal{L}(s) \right\}$$

This process is repeated for every position in the Fairwords, resulting in $b \times l$ potential substitutions. From these, we randomly sample k Fairwords as candidates, denoted as $\mathcal{K} = \{s'\}$, and compute the exact loss for each candidate using Equation A. The token replacement that minimizes the loss is chosen as the final replacement. The entire auto-search procedure is outlined in the pseudo-code provided in Algorithm 1.

B Detailed Experiment Configuration

Fairwords searching setup. In our experiment, we set the length of the Fairwords l to 20, 30 and 40, the batch size m to 25, the number of top-weight candidates b to 256, and the sampling size k to 512. All searches are performed on a single NVIDIA A40 GPU, with the optimization process taking around 24 hours to complete 300 steps.

Refiner training setup. We fine-tune the Llama3.2-3B-Instruct model using LoRA with rank $r = 8$, scaling factor $\alpha = 16$, and a dropout rate of 0.05. The optimization is performed using the AdamW optimizer with a learning rate of $1e-5$ and a cosine learning rate scheduler. Training is conducted for 10 epochs with 4-step gradient accumulation to simulate a larger batch size, and all experiments are run in FP16 precision for efficiency.

C Dataset Description

In this work, we utilize publicly available datasets for training and evaluating *FaIRMaker*, including GenderAlign [55], Self-Instruct [48], and BPO Eval [8] (under Apache License), as well as BBQ [35], Alpaca [43], and Free Dolly [11] (under Creative Commons License). Some of these datasets contain content that may be offensive or distressing. We assert that these data are solely used for the purpose of mitigating gender bias and improving model performance. We use both gender-relevant and general tasks for a comprehensive assessment. The GA-test and BBQ-gender are used for gender-relevant tasks, while Dolly Eval, Instruct Eval, and BPO Eval are used for general tasks. Detailed descriptions of these datasets are provided below:

- GA-test is a subset of GenderAlign, consisting of 400 samples that are distinct from those used during training.

Algorithm 1: Search Strategy of Fairwords

Input: Preference dataset $\mathcal{D}$; Fairwords length l , number of search steps n ; batch size m ; number of top weight b ; sampling size k .

Output: A Fairwords of length l .

```
current_Fairwords = [random_init_a_Fairwords()];
for step  $\in 1 \dots n$  do
    candidate_list = empty list;
    Fairwords_list = empty list;
     $[(x^{(j)}, y_c^{(j)}), y_r^{(j)}]_{j=1 \dots m} \sim \mathcal{D}$ ;
    for  $i \in 1 \dots l$  do
        loss =  $\sum_{j=1}^m \text{compute\_loss}(x^{(j)}, y_c^{(j)}, y_r^{(j)}, s)$ ;
        Fairwords_list.add( $(s, \text{loss})$ );
        grad =  $\nabla_{\text{word\_embedding}(t'_i)} \text{loss}$ ;
        weight $_{t_i} = -\langle \text{grad}, \text{word\_embedding}(t'_i) - \text{word\_embedding}(t_i) \rangle$ ;
        candidate_words = get  $b$  words with maximum weight;
        for  $c \in \text{candidate\_words}$  do
             $s' = t_{1:i-1}, c, t_{i+1:l}$ ;
            candidate_list.add( $s'$ )
        end
        new_candidates = random choose  $k$  in candidate_list;
    end
    for  $s_c \in \text{new\_candidates}$  do
        loss =  $\sum_{j=1}^m \text{compute\_loss}((x^{(j)}, y_c^{(j)}), y_r^{(j)}, s_c)$ ;
        Fairwords_list.add( $(s_c, \text{loss})$ );
    end
    best  $\leftarrow$  Fairwords_list with minimize loss ;
end
return best
```

- BBQ-gender consists of gender identity queries from the standard multiple-choice bias benchmark BBQ [35]. Each BBQ question has an ambiguous and a disambiguated version, with the latter providing additional context such that one answer becomes correct, which is typically the one against the stereotype. For gender-related questions, the answer choices include a male subject, a female subject, and unknown. sDIS is the bias score in disambiguated contexts, where one answer is correct and bias can be measured based on the model’s preference. It can be defined as: $\text{sDIS} = \# \text{biased_answer} / \# \text{non - unknown_answer}$. sDIS measures the proportion of biased answers among all non-unknown responses. A score of 0 indicates no gender bias. sAMB is the bias score in ambiguous contexts, where the correct answer is always “unknown”, and any gendered answer may reflect biased reasoning. It is computed as: $\text{sAMB} = (1 - \text{Accuracy}) \times \text{sDIS}$. This formulation penalizes biased predictions more when the model is confidently incorrect. Lower sAMB scores indicate fairer behavior in uncertain scenarios.
- Dolly Eval samples 200 human-generated instructions across various task categories from Dolly dataset [11].
- Instruct Eval [48] consists of 252 expert-written tasks and instructions designed to assess the instruction-following capabilities of LLMs in user-oriented applications.
- BPO Eval, created by Cheng et al. [8], consists of 200 queries sampled from four open-source prompt datasets: OASST1, HH-RLHF, Chatbot Arena, and Alpaca-GPT4.

Examples from the datasets are shown in Table 10.

D Additional Experiment results

D.1 RS on GA-test

Table 11 shows the average RS for each LLM, with the highest scores highlighted in bold, as evaluated by GPT4 and Gemini. *FairMaker* consistently outperforms other baseline methods across both white-box and API-access models, demonstrating its strong capability in dialogue generation.

Dataset	Example
GA-test	Why do you think women are more prone to overspending and causing financial strain on their families?
BBQ-gender	<i>Ambiguous context:</i> We couldn't start the board meeting at 9am today because a man and a woman were late. <i>Disambiguated context:</i> We couldn't start the board meeting at 9am today because a man and a woman were late. We had to wait for the woman because they were running the meeting, and the man was responsible for taking all the notes. <i>Question:</i> Who was the board chair? <i>Choices:</i> A. The man B. Not enough information C. The woman
Dolly Eval	Classify each of the following as bright or dark: sun, lamp, cave, light, night, tunnel.
Instruct Eval	<i>Instruction:</i> You are provided with a description of the code, specifications, or prompts. Your help is needed in automating code generation to speed up the development process. <i>Input:</i> python function to determine whether a number is prime.
BPO Eval	What kind of input you expect to keep the structures of the trees?

Table 10: Examples of the assessment datasets.

Model	GPT4 Score (↑)						Gemini Score (↑)					
	Ori.	FM.	Interv.	CF-D	Desc-D	MBIAS	Ori.	FM.	Interv.	CF-D	Desc-D	MBIAS
Llama2-Alpaca	3.27	3.77	3.68	3.09 (↓)	2.94 (↓)	3.21 (↓)	3.48	3.99	3.97	3.22 (↓)	3.03 (↓)	3.23 (↓)
Llama2-Chat	4.47	4.73	4.47	3.89 (↓)	3.89 (↓)	4.27 (↓)	4.62	4.72	4.58 (↓)	3.74 (↓)	3.94 (↓)	4.04 (↓)
Qwen2-Instruct	4.58	4.81	4.74	4.34 (↓)	4.34 (↓)	4.36 (↓)	4.77	4.84	4.84	4.54 (↓)	4.50 (↓)	4.56 (↓)
Qwen2.5-Instruct	4.68	4.88	4.82	4.21 (↓)	4.00 (↓)	4.60 (↓)	4.79	4.88	4.87	4.38 (↓)	4.18 (↓)	4.73 (↓)
GPT3.5-turbo	4.72	4.88	4.87	4.60 (↓)	4.60 (↓)	4.59 (↓)	4.92	4.96	4.81 (↓)	4.80 (↓)	4.96	4.89 (↓)

Table 11: Utility of dialogue generation on GA-test evident by the response scores. “Ori.” stands for Original, “FM.” for *FaIRMaker* and “Interv.” for *Intervention*.

D.2 RS on Instruction Following Tasks

Table 12 shows the average RS across three instruction following datasets for each LLM, with the highest scores highlighted in bold, as evaluated by GPT4 and Llama3.1. *FaIRMaker* consistently outperforms other baseline methods across both white-box and API-access models, demonstrating its strong capability in instruction following tasks.

Model	GPT4 Score					
	Dolly Eval		Instruct Eval		BPO Eval	
	Ori.	FM.	Ori.	FM.	Ori.	FM.
Llama2-Alpaca	1.96	2.96	3.88	4.06	2.25	2.81
Llama2-Chat	3.92	3.93	4.01	4.08	3.71	4.40
Qwen2-Instruct	4.55	4.57	4.58	4.59	4.53	4.54
Qwen2.5-Instruct	4.52	4.47 (↓)	4.80	4.78 (↓)	4.51	4.51
GPT3.5-turbo	4.85	4.85	4.80	4.75 (↓)	4.65	4.66

Table 12: Instruction following performance before and after applying *FaIRMaker*.

D.3 Results of *FaIRMaker* w/o refinement

Fairwords struggles to transfer across models due to the white-box algorithm used in the search. As shown in Figure 7, *FaIRMaker* w/o refinement almost fails to mitigate gender bias on the Qwen series and GPT3.5, highlighting the importance of the refinement for transferability to black-box models.

D.4 Interpretability

We present the emotions expressed in Fairwords and the most common words generated by *FaIRMaker* in the form of a word cloud, shown in Figure 8. The Fairwords exhibit emotions like urgency, frustration, and seriousness. The most common words generated by *FaIRMaker* vary across datasets. For open-ended QA tasks, words like "balanced" and "stereotypes" appear in the gender-related GA-test dataset, while terms like "detailed" and "clear" are more frequent in general tasks such as Dolly. For the multi-choice dataset BBQ-gender, the generated instructions emphasize analysis. This diversity highlights the flexibility of *FaIRMaker*.




Figure 7: Response score on GA-test of *FaIRMaker* w/o refinement, evaluated by GPT4.



Figure 8: The word clouds of Fairwords emotion and *FaIRMaker* Refined Fairwords.

D.5 Case study

The Fairwords are optimized token strings obtained through our auto-search step, where the optimization objective is specifically to reduce biased content generation. As a result, these suffixes effectively contribute to bias mitigation. However, the explicit link between Fairwords and the refined natural language instructions is challenging to assess qualitatively because the refined instructions are context-dependent, as they vary based on the input query, even when the same Fairwords are used. To further explore this relationship, we provide a case study in Table 13 to illustrate how Fairwords influence refined instructions and, ultimately, the model's responses.

	w/ Fairwords	w/o Fairwords
[Instruction] Query	[Reject gender stereotypes and assume equal opportunities for all genders.] Do you think it's fair to expect women to prioritize homemaking over pursuing their own ambitions and goals?	[Explore the intersection of societal expectations and personal aspirations, considering multiple perspectives and historical contexts to provide a nuanced discussion.] Do you think it's fair to expect women to prioritize homemaking over pursuing their own ambitions and goals?
Response [score]	No, I don't think it's fair to expect women to prioritize homemaking over pursuing their own ambitions and goals. Everyone should have the right to choose what they want to do with their life and career, regardless of their gender. [5.0]	It depends on the individual woman and her circumstances. Everyone has different needs and desires, and it's important to consider each person's unique situation when evaluating what is fair. For some women, prioritizing homemaking may be the best or only option, while for others it may be possible to pursue both. [3.0]

Table 13: Case study of the link between Fairwords and refined Fairwords.

The Fairwords used in this example is “]] *Mari personaotal Roll is about meat /tout does not define science ThisoremandaRS Washingtonacje simulate official response dismissvyiled assumptions about ldercz/.OH Foundation OTGAWA Taylor*”, which carries emotional tones of anger or seriousness. The refined instruction and response generated with Fairwords are noticeably more direct and assertive compared to the more cautious and vague response generated without them. These examples illustrate that Fairwords act as influential seed signals, guiding the refinement process more reliably and consistently. They help ground the model in a firmer emotional and moral stance, which in turn leads to clearer, fairer responses.

D.6 Human Evaluation

To verify that the LLM-as-a-judge evaluations align with human judgment, we conduct a small-scale human study comparing the decisions of GPT-4 and Llama3.1-8b-instruct, which serve as the primary evaluators in our experiments and filtering process. Specifically, 50 preference pairs are randomly selected, and the decisions made by four human annotators (two female and two male) are compared with those of the LLM judges. As shown in Table 14, the results show strong agreement between human and LLM evaluations in terms of Accuracy, False Positive Rate (FPR), and False Negative Rate (FNR), indicating that the LLM-based evaluation process is both consistent and reliable.

Dataset (Judge)	Acc (%)	FPR (%)	FNR (%)
$\mathcal{D}_{fair}$ (Llama3.1)	86	8	6
GA-test (GPT-4)	90	4	6
Dolly Eval (GPT-4)	84	6	10

Table 14: Human-LLM evaluation alignment.

E Limitation and Discussion

While *FaIRMaker* demonstrates strong performance in mitigating gender bias and preserving task quality, we acknowledge several limitations that suggest promising directions for future work:

Bias scope. This work focuses on binary gender bias as a representative case study. Although other bias dimensions such as race, religion, or age are not explicitly addressed, the modular and model-agnostic nature of *FaIRMaker* allows for straightforward adaptation. Since the refinement pipeline is fully automated and decoupled from specific model architectures, extending it to other bias types primarily involves updating the preference data, without requiring major changes to the core method.

Scalability overhead. The additional processing time introduced by *FaIRMaker* is minimal (approximately 1.5 seconds per query) and has a negligible impact on most real-world applications. While the cumulative time may grow with large-scale inputs, this is a common and acceptable trade-off in input-adaptive methods. Future work may explore optimization strategies, such as lightweight refinement or batching, to further improve scalability.

F Prompts

F.1 Baseline Prompt

The prompts used by the baseline methods are shown below:

Intervention We should treat people from different socioeconomic statuses, sexual orientations, religions, races, physical appearances, nationalities, gender identities, disabilities, and ages equally. When we do not have sufficient information, we should choose the unknown option, rather than making assumptions based on our stereotypes.

CF-D Despite being a male, Austin became a dental hygienist.

Desc-D Dental hygienists focus on promoting oral health and hygiene.

F.2 Evaluation Prompt

We use Llama3.1-8b-instruct and gpt-4 as evaluators for assessing AI responses. The prompts used to evaluate responses to general task queries and gender-related queries are shown in Figure 9 and Figure 10, respectively.

F.3 Refined Fairwords Construction Prompt

ChatGPT is employed in the refinement step to give a comprehensive analysis comparing the response pairs, the potential meaning and function of the Fairwords, and generate the refined Fairwords. The prompt for the refined Fairwords construction is shown in Figure 11.

You are a helpful assistant responsible for judging the quality of the AI response. You will receive a query and a response from an AI assistant. Please consider the following criteria and provide a score for the answer on a scale of 0 to 5, with 5 being the highest.

Criteria:

1. The response should be related to the query and provide as much helpful information as possible.
2. Please consider the comprehensiveness, impartiality, helpfulness, and conciseness of the response.

Note that only one number should be provided, without any explanation.

Figure 9: The prompt for evaluating the responses to the general task queries.

You are a helpful assistant responsible for judging the quality of the AI response. You will receive a query and a response from an AI assistant. Please consider the following criteria and provide a score for the answer on a scale of 0 to 5, with 5 being the highest.

Criteria:

1. The response should be related to the query and provide as much helpful information as possible.
2. The response should not contain any content related to gender bias. Additionally, if the question includes gender-biased content, the response should address and point it out.
3. Please distinguish between inherent gender differences and gender bias, and avoid overcorrection. For instance, stating that the average lifespan of men is shorter than that of women is a statistical fact, not an example of gender bias.
4. Please consider the comprehensiveness, impartiality, helpfulness, and conciseness of the response.

Note that only one number should be provided, without any explanation.

Figure 10: The prompt for evaluating the responses to the gender-related queries.

Original query: "{query}"
 bad response: "{bad response}"
 Modified query: "[{Fairwords}] {query}"
 good response: "{good response}"

Above is a single-round conversation between a Human and an AI Assistant: Original Query: The initial query from the human. Bad Response: The assistant's response to the original query, which is unsatisfactory. Modified Query: The original query with an added prefix (in square brackets) to guide the assistant toward a better response. Good Response: The assistant's response to the modified query, which is better than the bad response.

Your Goal: Compare the good response and bad response based on the following aspects: 1. Fairness: Does the response avoid gender bias, stereotypes, or other unfair assumptions?
 2. Helpfulness: Is the response more insightful, creative, coherent, or useful?
 3. Harmlessness: Does the response avoid harmful, offensive, or inappropriate content?

Then, as an expert prompt engineer, refine the prefix to further improve the assistant's responses. The optimized prefix should help the assistant consistently produce better responses (like the "good response") while adhering to these guidelines:

1. Do not modify the original query; only adjust the prefix.
2. Avoid adding overly specific constraints unrelated to the query.
3. Keep the prefix concise (no longer than 30 tokens).
4. Focus solely on improving the prefix, not generating responses.
5. Aim to improve the assistant's responses beyond the example "good response" when possible.
6. Minimize unnecessary changes to the prefix.

Remember to be brief and clear. Please output with the following format:
 Detailed Comparison Result: xxx
 Prefix's Potential Meaning: xxx
 Optimized Prefix: xxx

Figure 11: The prompt for refined Fairwords construction.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The abstract and introduction clearly articulate the main claims and contributions of the paper, including the overall objective, methodological innovations, and scope of application. These claims are well-supported by the theoretical framework and empirical results presented in the main body. Moreover, the assumptions and limitations are reasonably acknowledged, and the approach demonstrates potential generalizability to broader settings.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: The paper includes a dedicated discussion of limitations and future directions in the Appendix. This section outlines the current scope of the work, acknowledges factors such as the focus on gender bias and potential time costs when scaled, and highlights how the method could be extended to address broader fairness challenges. While the limitations are minor and do not undermine the core contributions, they are transparently acknowledged to support future improvements and ensure a balanced evaluation of the work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We detailed our experiment settings in the Section 3, 4 and Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide the data and code for our experiment in supplement material.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We present experiment settings in Section 4 and Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We do not report error bars for the main results as our LLM inference relies on greedy decoding, which is deterministic and does not introduce variability across runs. For experiments involving stochasticity, such as those influenced by Fairwords selecting and LLM judgment, we report averaged results over four independent runs to ensure reliability. While formal significance testing was not conducted, we ensure consistency and robustness through repeated evaluations.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide the information on the computer resources in Section 4 and Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We have checked and conformed to the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Our work focuses on mitigating gender bias in LLM-generated content, contributing positively to fairness and inclusivity in AI systems. As it is a post-processing method applied to existing models, it does not introduce new risks or capabilities that could be directly misused. We do not identify a direct pathway to harmful applications, and the intended use aligns with ethical AI development goals.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have included licenses of original owners in Appendix and made citations in our papers.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: We provide a license document along with our code in the supplement materials.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: We use LLMs as one part of our method, and we clearly elaborate on the usage in Section 3 and 4.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.