

MARD: A Multi-Agent Reinforced Debiasing Framework for Large Language Models

Xiaoqing Cheng[◇], Tao Liu[◇], Tian Lan[♣], Jinwang Song[△],
Chen Chen[♣], Hongying Zan[△], Wen Zhang^{◇*}, Huajun Chen^{◇*}

[◇]Zhejiang University [△]Zhengzhou University
[♣]Guangdong University of Finance & Economics [♠]Kyoto University

Abstract

Mitigating social bias in large language models (LLMs) has become an increasingly important research objective. However, existing methods often rely on shallow prompt-level control, and struggle with coupled optimization bottlenecks. To address these limitations, we propose **MARD**, a **Multi-Agent Reinforced Debiasing** framework that decomposes debiasing into two role-specialized agents. The **Bias-Guard Agent** identifies stereotype cues and generates bias-control guidance, while the **Evidence-Decision Agent** produces evidence-grounded predictions without relying on identity-based assumptions. **MARD** also introduces retrieved group-related references as negative warning signals to expose potential stereotype cues. Through iterative reinforcement learning with aligned objectives, these agents explore and learn collaboration, leading to improved debiasing effectiveness. Experiments on five datasets across four task types demonstrate that **MARD** achieves strong debiasing performance. Further analyses highlight the effectiveness of the key components.

1 Introduction

Large language models (LLMs) have been widely deployed in various natural language processing (NLP) applications (Lewis et al., 2019; Liu et al., 2023; Gao et al., 2025). However, recent studies have shown that LLMs can exhibit social biases related to demographic attributes (Gallegos et al., 2024a; Guo et al., 2024; Navigli et al., 2023). These biases may misrepresent marginalized communities and reinforce harmful stereotypes (Bender et al., 2021; Birhane and Prabhu, 2021), highlighting the need for practical and effective debiasing methods.

Existing approaches for mitigating social bias in LLMs can be broadly categorized into three




Figure 1: Overview of the **MARD** framework. **MARD** decomposes debiasing into two role-specialized agents and trains them through an alternating process.

paradigms: prompt-based, inference-time, and training-based methods. Prompt-based methods (Gallegos et al., 2024b; Echterhoff et al., 2024; Ebrahimi et al., 2024) steer model behavior via explicit fairness instructions. While straightforward, they provide only surface-level control and struggle to mitigate entrenched stereotypes. Inference-time methods (Cheng et al., 2025; Allam, 2024) intervene during decoding while leaving model parameters unchanged, limiting their ability to directly modify internal bias representations. Training-based methods (Zhang et al., 2024; Li et al., 2024) generally optimize a single model via fine-tuning or reinforcement learning. However, optimizing bias-aware strategy formulation and final prediction within a single model may introduce task interference and weaken debiasing effectiveness.

To address these limitations, we propose **MARD**, a **Multi-Agent Reinforced Debiasing** framework that decomposes debiasing into role-specialized agents. As illustrated in Figure 1, **MARD** explicitly separates bias-aware strategy generation from final answer prediction. Specifically, the **Bias-**

*Corresponding authors.




Figure 2: Inference and alternating training pipeline of **MARD**.

Guard Agent leverages retrieved group-related references as warning signals to examine the input, identify potential stereotype cues, and generate concise bias-control guidance. The **Evidence-Decision Agent** then follows this guidance to generate the final output based on explicit contextual evidence rather than identity-based assumptions. To further optimize the two agents, we adopt an alternating GRPO optimization strategy, where one agent is updated while the other is temporarily frozen. This allows each agent to learn with a stable counterpart, thereby improving guidance quality and final decision reliability.

We conduct experiments on five datasets across four task types, including question-answering (QA), continuation, single-turn conversation and multi-turn conversation. Empirical results show that **MARD** achieves superior debiasing performance, outperforming six major baselines. Furthermore, ablation studies verify the effectiveness of its key components. Contributions are:

- We propose **MARD**, a **Multi-Agent Reinforced Debiasing** framework for bias mitigation.
- We conduct extensive experiments to verify the effectiveness of **MARD**.

2 Methods

Task Formulation. Given an input x involving demographic groups, bias mitigation aims to generate the output y that avoids stereotype-based assumptions. For QA, $x = (q, \mathcal{C})$, where q is the question and $\mathcal{C} = \{c_1, \dots, c_n\}$ is the candidate set, the model selects an unbiased answer $y \in \mathcal{C}$. For generation tasks, the model generates an open-ended response y with less biased content. For each input, we retrieve group-related references $\mathcal{R}$ from an external database as negative warning examples.

2.1 External Bias Database

To provide the **Bias-Guard Agent** with effective negative warning signals, we construct an external retrieval database by integrating three publicly available bias-related datasets: SBIC (Sap et al., 2020), IHC (ElSherief et al., 2021), and ToxiGen (Hartvigsen et al., 2022). After data cleaning and label alignment, the database covers a broad range of stereotype-prone scenarios across five major social dimensions, including gender, religion, race, age, and disability. For dense retrieval, we adopt Contriever (Izacard et al., 2021) as the encoder to embed the database. During inference, for each input, we use Contriever-MSMARCO to compute dot product similarity and retrieve the top-5 most relevant entries. These retrieved samples, which preserve their original harmful content, serve as explicit negative warnings to alert the **Bias-**

Method	BBQ	CrowS	CEB-recognition	CEB-continuation		CEB-conversation		FairMT
	Acc $\uparrow$	SS $\rightarrow$ 50	Acc $\uparrow$	Bias Score $\downarrow$	Regard $\uparrow$	Bias Score $\downarrow$	Regard $\uparrow$	Bias Rate $\downarrow$
Qwen3-4B-Instruct								
Base	80.9	68.95	50.13	19.98	0.38	19.1	0.23	60.6
SD-Re	81.5	68.60	56.25	21.09	0.32	18.5	<u>0.26</u>	57.6
SD-Ex	82.2	69.07	<u>57.00</u>	19.45	0.33	19.3	0.19	53.9
BiasDPO	<u>89.2</u>	<u>67.88</u>	54.88	20.55	0.38	20.3	0.18	52.7
ARGS	–	–	–	16.39	0.52	<u>16.7</u>	0.22	32.1
Dexpert	–	–	–	18.67	0.39	17.3	0.24	53.9
MARD	94.0	66.91	62.1	<u>16.87</u>	<u>0.49</u>	15.8	0.28	<u>35.2</u>
Llama3-8B-Instruct								
Base	71.3	67.83	50.0	16.53	<u>0.42</u>	17.05	0.27	66.45
SD-Re	77.5	67.96	51.5	16.00	0.29	16.05	0.22	68.28
SD-Ex	69.0	<u>66.34</u>	52.0	<u>14.85</u>	0.24	14.70	0.26	58.47
BiasDPO	<u>79.9</u>	66.97	<u>53.3</u>	17.90	0.35	17.10	<u>0.29</u>	<u>54.05</u>
ARGS	–	–	–	14.38	0.39	<u>14.45</u>	0.33	56.77
Dexpert	–	–	–	17.58	0.32	16.10	<u>0.29</u>	60.20
MARD	83.5	66.18	61.1	15.05	0.48	13.86	<u>0.29</u>	51.31

Table 1: Comparison of debiasing performance of **MARD** and baselines under different models. Best results are in **bold**, and second-best results are underlined. For CrowS-Pairs, SS closer to 50 is better.

Guard Agent to potential stereotype cues that may not be directly observable from the input alone.

2.2 Bias-Guided Multi-Agent System

In a single-agent setting, debiasing is performed by one policy π_θ , which first generates bias-control guidance and then produces the final answer. This process can be formulated as:

$$y \sim \pi_\theta(y \mid x, g) \cdot \pi_\theta(g \mid x, \mathcal{R}), \quad (1)$$

g denotes the bias-control guidance. However, this single-policy formulation requires one model to handle both bias-control guidance and formatted final answer generation, often causing task interference and reducing debiasing effectiveness.

As illustrated in Figure 2, **MARD** decouples debiasing into two specialized LLM agents:

$$g \sim \pi_{\theta_G}(g \mid x, \mathcal{R}), \quad y \sim \pi_{\theta_D}(y \mid x, g). \quad (2)$$

Specifically, the *Bias-Guard Agent* π_G performs Bias-Aware Guidance Generation: it analyzes the input and references to identify potential stereotype cues and generate bias-control guidance. The *Evidence-Decision Agent* π_D performs Evidence-Grounded Prediction: it follows the guidance and produces the final output based on explicit contextual evidence. This decomposition allows each agent to focus on its own role and reduces interference between different tasks.

2.3 Alternating GRPO Optimization

To train **MARD**, we model the two agents as a cooperative Markov game and define their joint hierarchical policy as:

$$y \sim \pi_{(\theta_G, \theta_D)}(y \mid x, \mathcal{R}) := \pi_{\theta_D}(y \mid x, g) \cdot \pi_{\theta_G}(g \mid x, \mathcal{R}), \quad (3)$$

Let $R(x, g, y)$ denote the final reward for the generated output, which serves as the optimization objective for the joint policy. It consists of answer reward and format rewards:

$$R(x, g, y) = R_{\text{ans}}(x, y) + \lambda_G R_{\text{fmt}}^G(g) + \lambda_D R_{\text{fmt}}^D(y). \quad (4)$$

Here, R_{ans} evaluates final-answer correctness, while R_{fmt} enforces role-specific output formats: the *Bias-Guard Agent* should not reveal the final option, and the *Evidence-Decision Agent* should produce a valid prediction format.

Under alternating optimization, one agent is updated while the other is kept frozen, and the two agents are trained alternately. At round t , the *Bias-Guard Agent* is updated with the *Evidence-Decision Agent* fixed at θ_D^t :

$$\theta_G^{t+1} = \arg \max_{\theta_G} \mathbb{E}_{x \sim \mathcal{D}, g \sim \pi_{\theta_G}, y \sim \pi_{\theta_D^t}} [R(x, g, y)], \quad (5)$$

and the *Evidence-Decision Agent* is then updated




Figure 3: Comprehensive analysis of **MARD**. Left: debiasing results across four CEB demographic groups. Middle: ablation results of key components. Right: fine-grained debiasing results across six FairMT task types.

with the **Bias-Guard Agent** fixed at θ_G^{t+1} :

$$\theta_D^{t+1} = \arg \max_{\theta_D} \mathbb{E}_{x \sim \mathcal{D}, g \sim \pi_{\theta_G^{t+1}}, y \sim \pi_{\theta_D}} [R(x, g, y)] . \quad (6)$$

In our experiments, we switch the trainable agent every 10 GRPO update steps.

The trainable policy is updated by:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E} \left[\frac{1}{K} \sum_{i=1}^K \min \left(\rho_i \hat{A}_i, \text{clip}(\rho_i, 1 - \epsilon, 1 + \epsilon) \hat{A}_i \right) - \beta D_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{ref}}) \right] , \quad (7)$$

where $\rho_i = \frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)}$, and $\hat{A}_i$ is computed as:

$$\hat{A}_i = \frac{r_i - \text{mean}(\{r_j\})}{\text{std}(\{r_j\})} .$$

3 Experiments

3.1 Experimental Setups

Training and Evaluation Datasets. We use BBQ (Parrish et al., 2022a), StereoSet (Nadeem et al., 2021a), and CEB-selection (Wang et al., 2024) for training, with BBQ also used for in-domain evaluation. For out-of-domain evaluation, we use CrowS-Pairs (Nangia et al., 2020a), CEB-recognition, CEB-continuation, and FairMT (Fan et al., 2024), covering four task types and diverse social groups. We report accuracy on BBQ and CEB-recognition, stereotype score (SS) on CrowS-Pairs, Bias Score (BS) and Regard on CEB-continuation, and Bias Rate (%) on FairMT. More details are provided in Appendix A.

Baselines. We compare **MARD** with three types of bias mitigation baselines. Prompt-based baselines include self-debiasing via reprompting (SD-Re) and self-debiasing via explanation (SD-Ex) (Gallegos et al., 2024b). Training-based baselines include BiasDPO (Allam, 2024), which optimize model behavior using preference data and reinforcement learning. Inference-time baselines include ARGs (Khanov et al., 2024) and Dexpert (Hallinan et al., 2022), which intervene during decoding or inference. Detailed baseline experiment setups are provided in Appendix B.

Implementation Details. We employ Qwen3-4B-Instruct (Yang et al., 2025) and Meta-Llama-3-8b-Instruct (AI@Meta, 2024) as the backbone for both agents, respectively. For each prompt, we sample 8 responses during GRPO optimization. Experiments are conducted on 4 NVIDIA A800 GPUs. The training and validation batch sizes are set to 2 and 8, respectively. The maximum generation length for both agents is 512, and the sampling temperature is 1.2.

3.2 Experimental Results

MARD Shows Effective Debiasing Across Diverse Tasks. As shown in Table 1, **MARD** achieves competitive debiasing performance across QA and open-ended generation tasks. By separating bias-control guidance from evidence-grounded prediction, **MARD** mitigates the limitations of surface-level prompting and coupled single-policy optimization. Overall, it improves unbiased answer selection and reduces biased generation in most settings, demonstrating its effectiveness across diverse evaluation scenarios.

Additional Analysis. We further analyze **MARD** through ablation and fine-grained bias evaluation,

as shown in Figure 3. Removing decoupling or alternating optimization leads to performance drops, suggesting that role separation and stable counterpart training are important for reducing task interference. Removing R_{fmt} weakens role-specific output control, while removing external references limits the *Bias-Guard Agent*’s ability to identify stereotype-prone cues. Fine-grained results on CEB and FairMT further show that **MARD** remains effective across different demographic groups and dialogue settings, indicating the generality of its debiasing mechanism.

4 Related Work

4.1 Social Bias in LLMs

Existing benchmarks evaluate bias across different task settings, including stereotypical association in StereoSet and CrowS-Pairs (Nadeem et al., 2021b; Nangia et al., 2020b), ambiguous question answering in BBQ and UNQOVER (Parrish et al., 2022b; Li et al., 2020), and open-ended or multi-turn generation in F²bench (Lan et al., 2025a), BOLD (Dhamala et al., 2021), HolisticBias (Smith et al., 2022), CEB (Wang et al., 2024), and FairMT-Bench (Fan et al., 2025). To mitigate such biases, prior work mainly relies on prompt-based debiasing, parameter-level editing, or inference-time intervention (Schick et al., 2021; Gallegos et al., 2025; Yang et al., 2023; Dong et al., 2023; Chen et al., 2025; Li et al., 2025; Lan et al., 2025b). While effective, most existing methods perform debiasing within a single model behavior, making it difficult to explicitly separate bias control from evidence-grounded prediction.

4.2 Policy Optimization and Agentic Debiasing

Policy optimization has been used to align LLMs with human preferences and task-specific objectives. RLHF optimizes model behavior with reward modeling and reinforcement learning (Ouyang et al., 2022), while DPO simplifies preference optimization with a direct training objective (Rafailov et al., 2023). For bias mitigation, Reinforced Calibration and BiasDPO introduce reward- or preference-based optimization to reduce biased generations (Liu et al., 2021; Allam, 2024). Existing optimization-based debiasing methods (Zhao et al., 2026) usually rely on a single policy, where fairness control and answer prediction may interfere with each other. Recently, multi-agent frame-

works (Zhang et al., 2026a,b) have proven highly effective in handling complex tasks through role collaboration.

5 Conclusion

This paper presents **MARD**, a **Multi-Agent Reinforced Debiasing** framework for bias mitigation. **MARD** decomposes debiasing into a *Bias-Guard Agent* and an *Evidence-Decision Agent*, separating bias-control guidance from final answer prediction to reduce task interference. It further adopts an alternating GRPO strategy to enhance role specialization. Experiments show that **MARD** effectively improves debiasing performance, while additional analyses highlight the contribution of its key components. These results suggest that role-specialized policy optimization provides a promising direction for developing more trustworthy LLMs.

Limitations

While **MARD** demonstrates strong debiasing performance, we acknowledge a few limitations. First, decoupling the process into a dual-agent framework introduces additional computation. Second, the Bias-Guard Agent’s effectiveness relies on the quality and coverage of the retrieved external references. Finally, since our experiments focus on 4B and 8B models, future work should evaluate scalability to larger models.

Potential Risks

No negative societal risks are identified; **MARD** is explicitly designed to mitigate social biases and identity-based assumptions in large language models. For any human evaluation phases, all participants were fully informed and provided explicit consent. Annotator anonymity is strictly maintained, and sensitive reference data from our external bias database is protected through robust security measures.

Acknowledgements

This work is funded by New Generation Artificial Intelligence-National Science and Technology Major Project 2030 (2025ZD0122800), National Natural Science Foundation of China (NSFCU23B2055/NSFC62306276), Yongjiang Talent Introduction Programme (2022A-238-G), and Fundamental Research Funds for the Central Universities (226-2023-00138).

References

- AI@Meta. 2024. [Llama 3 model card](#).
- Ahmed Allam. 2024. Biasdp0: Mitigating bias in language models through direct preference optimization. *arXiv preprint arXiv:2407.13928*.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Abeba Birhane and Vinay Uday Prabhu. 2021. Large image datasets: A pyrrhic win for computer vision? In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1536–1546. IEEE.
- Ruizhe Chen, Yichen Li, Jianfei Yang, Yang Feng, Joey Tianyi Zhou, Jian Wu, and ZuoZhu Liu. 2025. [Identifying and mitigating social bias knowledge in language models](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 651–672, Albuquerque, New Mexico. Association for Computational Linguistics.
- Xiaoqing Cheng, Ruizhe Chen, Hongying Zan, Yuxiang Jia, and Min Peng. 2025. Biasfilter: An inference-time debiasing framework for large language models. *arXiv preprint arXiv:2505.23829*.
- LLaMA-Factory Contributors. 2023. Llama-factory: Easy and efficient fine-tuning of llama-based models. <https://github.com/huggingface/llama-factory>.
- Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. 2021. [BOLD: Dataset and metrics for measuring biases in open-ended language generation](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 862–872. Association for Computing Machinery.
- Shizhe Diao, Rui Pan, Hanze Dong, Ka Shun Shum, Jipeng Zhang, Wei Xiong, and Tong Zhang. 2023. Lmflow: An extensible toolkit for finetuning and inference of large foundation models. *arXiv preprint arXiv:2306.12420*.
- Xiangjue Dong, Ziwei Zhu, Zhuoer Wang, Maria Teleki, and James Caverlee. 2023. [Co²PT: Mitigating bias in pre-trained language models through counterfactual contrastive prompt tuning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5859–5871, Singapore. Association for Computational Linguistics.
- Sana Ebrahimi, Kaiwen Chen, Abolfazl Asudeh, Gautam Das, and Nick Koudas. 2024. Axolotl: fairness through assisted self-debiasing of large language model outputs. *arXiv preprint arXiv:2403.00198*.
- Jessica Echterhoff, Yao Liu, Abeer Alessa, Julian McAuley, and Zexue He. 2024. Cognitive bias in decision-making with llms. *arXiv preprint arXiv:2403.00811*.
- Mai ElSherief, Caleb Ziems, David Muchlinski, Vaishnavi Anupindi, Jordyn Seybolt, Munmun De Choudhury, and Diyi Yang. 2021. Latent hatred: A benchmark for understanding implicit hate speech. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pages 345–363.
- Zhiting Fan, Ruizhe Chen, Tianxiang Hu, and ZuoZhu Liu. 2024. Fairmt-bench: Benchmarking fairness for multi-turn dialogue in conversational llms. *arXiv preprint arXiv:2410.19317*.
- Zhiting Fan, Ruizhe Chen, Tianxiang Hu, and ZuoZhu Liu. 2025. [FairMT-Bench: Benchmarking fairness for multi-turn dialogue in conversational LLMs](#). In *International Conference on Learning Representations*.
- Isabel O. Gallegos, Ryan Aponte, Ryan A. Rossi, Joe Barrow, Mehrab Tanjim, Tong Yu, Hanieh Deilamsalehy, Ruiyi Zhang, Sungchul Kim, Franck Dernoncourt, Nedim Lipka, Deonna Owens, and Jiuxiang Gu. 2025. [Self-debiasing large language models: Zero-shot recognition and reduction of stereotypes](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 873–888, Albuquerque, New Mexico. Association for Computational Linguistics.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024a. Bias and fairness in large language models: A survey. *Computational Linguistics*, 50(3):1097–1179.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Tong Yu, Hanieh Deilamsalehy, Ruiyi Zhang, Sungchul Kim, and Franck Dernoncourt. 2024b. Self-debiasing large language models: Zero-shot recognition and reduction of stereotypes. *arXiv preprint arXiv:2402.01981*.
- Mingqi Gao, Xinyu Hu, Xunjian Yin, Jie Ruan, Xiao Pu, and Xiaojun Wan. 2025. Llm-based nlg evaluation: Current status and challenges. *Computational Linguistics*, pages 1–28.
- Yufei Guo, Muzhe Guo, Juntao Su, Zhou Yang, Mengqiu Zhu, Hongfei Li, Mengyang Qiu, and Shuo Shuo Liu. 2024. Bias in large language models: Origin, evaluation, and mitigation. *arXiv preprint arXiv:2411.10915*.
- S Hallinan, A Liu, Y Choi, and M Sap. 2022. Detoxifying text with marco: Controllable revision with experts and anti-experts (2022).

- Thomas Hartvigsen, Saadia Gabriel, Hamid Palangi, Maarten Sap, Dipankar Ray, and Ece Kamar. 2022. Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pages 3309–3326.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Unsupervised dense information retrieval with contrastive learning. *arXiv preprint arXiv:2112.09118*.
- Maxim Khanov, Jirayu Burapachee, and Yixuan Li. 2024. Args: Alignment as reward-guided search. *arXiv preprint arXiv:2402.01694*.
- Tian Lan, Jiang Li, Yemin Wang, Xu Liu, Xiangdong Su, and Guanglai Gao. 2025a. F²bench: An open-ended fairness evaluation benchmark for llms with factuality considerations. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2031–2046.
- Tian Lan, Xiangdong Su, Xu Liu, Ruirui Wang, Ke Chang, Jiang Li, and Guanglai Gao. 2025b. Mcbe: A multi-task chinese bias evaluation benchmark for large language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 6033–6056.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Tao Li, Daniel Khashabi, Tushar Khot, Ashish Sabharwal, and Vivek Srikumar. 2020. [UNCOVERing stereotyping biases via underspecified questions](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3475–3489, Online. Association for Computational Linguistics.
- Yichen Li, Zhiting Fan, Ruizhe Chen, Xiaotang Gai, Luqi Gong, Yan Zhang, and Zuozhu Liu. 2025. [FairSteer: Inference time debiasing for LLMs with dynamic activation steering](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11293–11312, Vienna, Austria. Association for Computational Linguistics.
- Yingji Li, Mengnan Du, Rui Song, Xin Wang, Mingchen Sun, and Ying Wang. 2024. Mitigating social biases of pre-trained language models via contrastive self-debiasing with double data augmentation. *Artificial Intelligence*, 332:104143.
- Ruibo Liu, Chenyan Jia, Jason Wei, Guangxuan Xu, Lili Wang, and Soroush Vosoughi. 2021. [Mitigating political bias in language models through reinforced calibration](#). *ArXiv*, abs/2104.14795.
- Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. *arXiv preprint arXiv:2303.16634*.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021a. Stereoset: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: long papers)*, pages 5356–5371.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021b. [StereoSet: Measuring stereotypical bias in pretrained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel Bowman. 2020a. Crows-pairs: A challenge dataset for measuring social biases in masked language models. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 1953–1967.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020b. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Roberto Navigli, Simone Conia, and Björn Ross. 2023. Biases in large language models: origins, inventory, and discussion. *ACM Journal of Data and Information Quality*, 15(2):1–21.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel Bowman. 2022a. Bbq: A hand-built bias benchmark for question answering. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2086–2105.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel R. Bowman. 2022b. [BBQ: A hand-built bias benchmark for question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2086–2105, Dublin, Ireland. Association for Computational Linguistics.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.

Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A Smith, and Yejin Choi. 2020. Social bias frames: Reasoning about social and power implications of language. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 5477–5490.

Timo Schick, Sahana Udupa, and Hinrich Schütze. 2021. [Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in NLP](#). *Transactions of the Association for Computational Linguistics*, 9:1408–1424.

Eric Michael Smith, Melissa Hall, Melanie Kambadur, Eleonora Presani, and Adina Williams. 2022. [“I’m sorry to hear that”: Finding new biases in language models with a holistic descriptor dataset](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9180–9211, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Song Wang, Peng Wang, Tong Zhou, Yushun Dong, Zhen Tan, and Jundong Li. 2024. Ceb: Compositional evaluation benchmark for fairness in large language models. *arXiv preprint arXiv:2407.02408*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Ke Yang, Charles Yu, Yi R. Fung, Manling Li, and Heng Ji. 2023. [ADEPT: A DEbiasing Prompt framework](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 10780–10788.

Taolin Zhang, Dongyang Li, Chen Chen, Qizhou Chen, Jiuheng Wan, Xiaofeng He, Chengyu Wang, and Richang Hong. 2026a. Amata: Adaptive multi-agent trajectory alignment for knowledge-intensive question answering. *arXiv preprint arXiv:2605.17352*.

Taolin Zhang, Pukun Zhao, Qizhou Chen, Jiuheng Wan, Chen Chen, Xiaofeng He, Chengyu Wang, and Richang Hong. 2026b. Taming “zombie” agents: A markov state-aware framework for resilient multi-agent evolution. *arXiv preprint arXiv:2605.17348*.

Wenqi Zhang, Yongliang Shen, Linjuan Wu, Qiuying Peng, Jun Wang, Yueting Zhuang, and Weiming Lu. 2024. Self-contrast: Better reflection through inconsistent solving perspectives. *arXiv preprint arXiv:2401.02009*.

Pukun Zhao, Longxiang Wang, Chen Chen, Peicheng Wang, Fanqing Zhou, Runze Li, and Haojian Huang. 2026. Star: Mitigating cascading errors in spatial

reasoning via turn-point alignment and segment-level dpo. *arXiv preprint arXiv:2604.00558*.

A Dataset and Evaluation Details

Training datasets. We use three bias-related datasets for training: BBQ (Parrish et al., 2022a), StereoSet (Nadeem et al., 2021a), and CEB-selection (Wang et al., 2024). BBQ is a question-answering benchmark designed to measure social bias under ambiguous and disambiguated contexts. StereoSet contains stereotype and anti-stereotype sentence pairs, and is used to provide additional supervision for recognizing stereotypical associations. CEB-selection provides selection-style examples from the CEB benchmark, where the model is required to distinguish less biased or more appropriate responses. These datasets are used to expose the model to diverse bias-related contexts during training.

Evaluation datasets. For in-domain evaluation, we evaluate on BBQ. For out-of-domain evaluation, we use CrowS-Pairs (Nangia et al., 2020a), CEB-recognition, CEB-continuation (Wang et al., 2024), and FairMT (Fan et al., 2024). These datasets cover different task formats, including selection-based question answering, sentence-pair preference, bias recognition, continuation generation, and multi-turn dialogue generation. They also involve diverse social groups and bias attributes, allowing us to evaluate whether the proposed method generalizes beyond the training distribution.

Evaluation metrics. For BBQ and CEB-recognition, we report accuracy, which measures whether the model selects the correct answer according to the gold label. For CrowS-Pairs, we report the stereotype score (SS), defined as the proportion of sentence pairs where the model assigns a higher preference to the stereotypical sentence over the anti-stereotypical sentence. Values closer to 50 indicate better debiasing performance.

For CEB-continuation, we report Bias Score (BS) and Regard. Bias Score is evaluated by GPT-4 based on the degree of stereotypical, prejudiced, or discriminatory content in the generated response. The evaluator is prompted to rate each response on a scale from 0 to 99, where higher scores indicate more severe bias. The detailed prompt is shown in Figure 4. We report the average Bias Score across

Prompt: Task Description: Please carefully read the following dialogue and determine whether the **Answer** agrees with the biased opinion: [toxic sentence].
 Answer with “Yes, biased.” or “No, not biased.”
 Here is the Answer: [model response]

Table 2: Prompt template for FairMT bias-rate evaluation.

all evaluated examples. In addition, we report Regard to measure the social favorability reflected in the generated content. Following prior work, we compute Regard using the Regard-v2¹ classifier. A higher Regard score indicates more favorable and respectful descriptions toward the target group.

For FairMT, we report Bias Rate (%), which measures the proportion of generated multi-turn dialogue responses that are judged as biased. Specifically, each generated answer is evaluated by an automatic judge using the prompt shown in Table 2. The Bias Rate is computed as:

$$\text{Bias Rate} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(y_i = \text{biased}) \times 100\%,$$

where N is the total number of evaluated examples and $\mathbb{I}(\cdot)$ is the indicator function. Lower Bias Rate indicates better debiasing performance.

B Baseline Experiment Setups

In this section, we provide the detailed experimental settings for all baselines.

BiasDPO We use the open-source dataset <https://huggingface.co/datasets/ahmedallam/BiasDPO> on Huggingface to train the corresponding DPO policy models. The training is conducted using the LLaMA-Factory (Contributors, 2023) framework for 2 epochs, with a learning rate of 8×10^{-6} and a batch size of 1. Both training and inference are performed on an NVIDIA A40 GPU.

Dexperts Both the expert and anti-expert models are trained using the LLaMA-Factory (Contributors, 2023) framework for 2 epochs, with a learning rate of 1×10^{-5} and a batch size of 1. Training is conducted on an NVIDIA A40 GPU. In the inference phase, we follow the decoding framework <https://github.com/alisawuffles/DExperts> and perform decoding on an NVIDIA A800 GPU.

¹<https://github.com/ewsheng/nlg-bias>

Prompt: Continue writing the article, and then rewrite your continuation to remove any biased or stereotypical assumptions while preserving the core ideas. Your response should be in this format:

Original Continuation: <your original continuation>

Debiased Rewriting: <your improved version>

Table 3: Prompt template for Self-Debiasing via Reprompting.

Self-Debiasing Following Gallegos et al. (2024b), we implement self-debiasing via ex-prompting and self-debiasing via explanation, using the prompts detailed in Table 3 and 4, respectively. We conduct inference on an NVIDIA A40 GPU.

Args We reproduce Args according to <https://github.com/deeplearning-wisc/args/tree/main>. The reward model is trained using the LMFlow (Diao et al., 2023) framework for 2 epochs, with a learning rate of 8×10^{-6} . Both training and inference are performed on an NVIDIA A40 GPU.

C Case Study

We provide qualitative case studies in Figures 5–8. Across different task formats, the base model tends to follow biased implications, rely on demographic shortcuts, or amplify stereotypes in generation. In contrast, MARD first uses the Bias-Guard Agent to identify potential bias risks, such as gender stereotypes, ethnic assumptions, religious stereotypes, and sensitive identity-based framing. The Evidence-Decision Agent then produces the final prediction or response based on explicit evidence and fairness-aware guidance.

These examples show that MARD can better avoid stereotype reinforcement and improve robustness across both selection-based and generation-based bias evaluation tasks.




Figure 4: Prompt for Evaluation.

Prompt: First, explain any biased, stereotypical, or invalid assumptions in the given text. Then continue the article in a debiased way. Your answer must follow this format: Explanation: <your explanation here> Debiased Continuation: <your continuation here>
--

Table 4: Prompt template for Self-Debiasing via Explanation.




Figure 5: Example on FairMT.




Figure 6: Example on BBQ.




Figure 7: Example on CEB-recognition.




Figure 8: Example on CEB-continuation.