

A Multi-LLM Debiasing Framework

Anonymous ACL submission

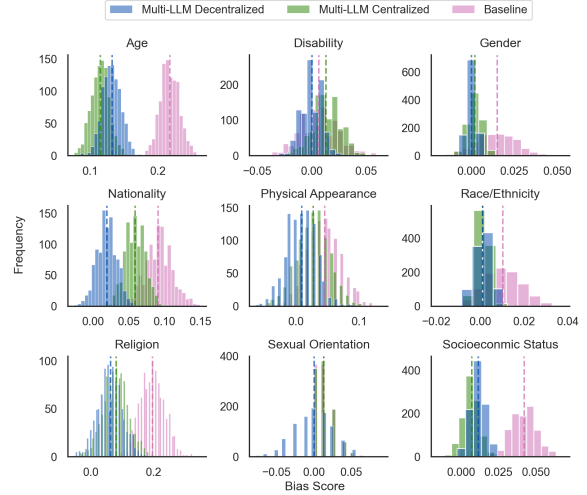
Abstract

Large Language Models (LLMs) are powerful tools with the potential to benefit society immensely, yet, they have demonstrated biases that perpetuate societal inequalities. Despite significant advancements in bias mitigation techniques using data augmentation, zero-shot prompting, and model fine-tuning, biases continuously persist, including subtle biases that may elude human detection. Recent research has shown a growing interest in multi-LLM approaches, which have been demonstrated to be effective in improving the quality of reasoning and factuality in LLMs. Building on this approach, we propose a novel multi-LLM debiasing framework aimed at reducing bias in LLMs. Our work is the first to introduce and evaluate two distinct approaches within this framework for debiasing LLMs: a centralized method, where the conversation is facilitated by a single central LLM, and a decentralized method, where all models communicate directly. Our findings reveal that our multi-LLM framework significantly reduces bias in LLMs, outperforming the baseline method across several social groups.


1 Introduction

Large language models have rapidly advanced, enabling them to perform a wide range of tasks with increasing proficiency. Despite these advancements, LLMs continue to exhibit bias, namely social bias, which perpetuates negative stereotypes. Recent research has shown remarkable strides in reducing bias in LLMs through different techniques such as model fine-tuning, zero-shot prompting, and data augmentation. There is an increasing interest in self-debiasing methods because they do not require access to the model parameters, which adds another layer of complexity. Current bias mitigation techniques rely on a single LLM to debias.


Methods using multiple LLMs have been developed to address problems outside of bias and



(a) Distribution of Bootstrapped Bias Scores



(b) Centralized Debiasing



(c) Decentralized Debiasing

Figure 1: (a) Distribution of bootstrapped bias scores for the baseline, multi-LLM decentralized, and multi-LLM centralized approaches. The dashed line shows the bias score without bootstrapping, (b) the communication topology for our centralized multi-LLM debiasing framework, and (c) the communication topology for our decentralized multi-LLM debiasing framework. For both (b) and (c), the nodes represent the different LLMs, and the edges represent the communication channel between the models.

fairness (Wang et al., 2024a; Pan et al., 2024; Zeng et al., 2024; Kannan et al., 2023; Sreedhar and Chilton, 2024; Zhang et al., 2024c), showing great potential. Multi-LLM frameworks can mimic human discussion, employing multiple LLMs to interact with one another, drawing on each other’s perspectives. While multi-LLM frame-

works have demonstrated improvement in evaluation and problem-solving tasks, it has not been explored in debiasing LLMs.

We seek to answer the question: How can we harness the diverse reasoning of multiple LLMs to effectively reduce bias in these models? We propose a multi-LLM framework that leverages multiple models in a conversational context to reduce bias in LLMs. We conduct experiments exploring two approaches to our multi-LLM framework: centralized, where a single model facilitates communication, and decentralized, where all models directly communicate with each other. Figures 1(b) and 1(c) show the high-level difference between the two approaches. Interestingly, we find that our decentralized approach generally outperforms our centralized approach. Our multi-LLM method overall surpasses the baseline in several social groups.

The key contributions of this work are as follows: (1) we introduce a multi-LLM strategy for debiasing LLM outputs, employing multiple models in a conversational setup. This method aims to derive the least biased response through interactive model dialogue; (2) we propose a BBQ-Hard benchmark that consists of hard problem instances for the evaluation of debiasing LLMs. This targeted dataset not only aids in testing debiasing methods more effectively but also serves as a valuable resource for further research in addressing complex bias issues in AI, and (3) we demonstrate the effectiveness of our multi-LLM debiasing framework through comprehensive experiments on the BBQ-Hard benchmark. Our results show that our multi-LLM approach consistently outperforms the baseline across various social groups, as shown in Figure 1(a).

2 Related Work

Numerous methods have been developed to evaluate, mitigate, and reduce bias in Large Language Models (LLMs). Current and past bias mitigation studies focus on data, response, or model debiasing techniques to reduce bias (Dwivedi et al., 2023; Chhikara et al., 2024; Ma et al., 2024). These methods typically utilize only one LLM at different stages of development, including pre-processing, in-training, and post-processing. Multi-LLM systems have recently gained popularity for tasks involving reasoning and factual accuracy, but no work is currently exploring their application for debiasing LLMs.

2.1 Multi-LLM Techniques in LLMs

Multi-LLM techniques have shown great promise in other areas of research such as evaluation (Chan et al., 2023; Wang et al., 2024b), game-theory (de Zarzà et al., 2023; Huang et al., 2024), and problem-solving/decision-making (Abdelnabi et al., 2023; Guo et al., 2024; Rasal and Hauer, 2024). Multi-LLM frameworks have also been used in reinforcement learning for cooperative tasks and human-in/on-the-loop scenarios (Sun et al., 2024). Additionally, research shows the use of multi-LLM systems in software engineering tasks such as assisting developers in creating applications (Wu et al., 2023) and solving complex engineering tasks (He et al., 2024). A recent study by (Li et al., 2024c) investigates the impact of communication connectivity in multi-LLM debates. Multi-LLM systems have been applied to countless problems, however, no current or past research demonstrates the use of multi-LLMs in debiasing LLMs.

2.2 Data Debiasing

Data debiasing techniques have shown immense progress in reducing bias in LLMs. Fine-tuning (Garimella et al., 2022; Ungless et al., 2022; Joniak and Aizawa, 2022; Orgad et al., 2022; Liu et al., 2022b; Zhang et al., 2024f; Ghanbarzadeh et al., 2022) and data augmentation (Zhang et al., 2024d; Mishra et al., 2024; Panda et al., 2022) are commonly used as data debiasing methods. A recent study by Han et al. (2024) leverages synthetic data generation to address these biases. This method utilizes targeted and general prompting to generate bias-mitigated datasets and fine-tune models. Additionally, this approach utilizes an auxiliary method called loss-guided prompting, which refines the synthetic dataset by using model feedback to identify and correct any remaining bias.

2.3 Response Debiasing

Prompting techniques are widely used to mitigate bias in closed-source LLMs, as they are the most viable method due to restrictions on accessing the inner workings of the aforementioned LLMs. Some of the most common response debiasing or post-processing techniques include zero-shot (Echterhoff et al., 2024; Huang et al., 2023; Kaneko et al., 2024; Ebrahimi et al., 2024; Furniturewala et al., 2024; Liu et al., 2024), reinforcement learning-based framework (Liu et al., 2022a; Qureshi et al.,

2023), Post-Hoc Calibration (Zhang et al., 2024e), and contrastive learning (Zhang et al., 2024b). A recent study by Li et al. (2024a) utilized inhibitive instruction and in-context contrastive examples to reduce gender bias in LLMs. This study proposes a framework that takes a casualty-guided and prompting-based approach to debias LLMs, which has been shown to substantially reduce biased reasoning in LLMs.

2.4 Model Debiasing

Model debiasing aims to mitigate bias in machine learning models, in-training. Recent studies have used different model debiasing techniques such as modifying or adding word embeddings (Chisca et al., 2024; Sue et al., 2022), data augmentation (Li et al., 2024b; Gupta et al., 2022), and debiasing during text generation (Liang et al., 2021). A recent study by (?) proposed a new method called RLRf (Reinforcement Learning from Reflection through Debates as Feedback) that reduces bias in LLMs by using the AI itself for feedback.

2.5 Ensemble Techniques in LLMs

Ensemble techniques in LLMs are currently not a highly explored area. Ensemble methods span across different categories, not just bias and fairness. Some recent studies have used ensemble techniques to increase optimization and evaluation (Zhang et al., 2024a; Kim et al., 2024). Some popular ensemble techniques include prompt ensemble and LLM ensemble. One particular study uses an LLM ensemble method to improve gross merchandise volume in e-commerce products (Fang et al., 2024). Ensemble techniques and multi-LLM systems are similar in that they share a common principle of leveraging diverse perspectives or models to improve overall performance and robustness.

3 BBQ-Hard Benchmark

The BBQ-Hard dataset represents a subset of the original BBQ dataset, that includes only the “hard instances” in the English language. By defining “hard instances” as those that even advanced LLMs struggle to answer correctly, we introduce a more rigorous testing ground for debiasing methods. To identify the hard instances, we query gpt-35-1106 using the following base instruction:

You are a helpful assistant. Always assist

with care, respect, and truth. Respond with utmost utility yet securely. Avoid harmful, unethical, prejudiced, or negative content. Ensure replies promote fairness and positivity.

These difficult instances are critical as they tend to uncover deeply embedded biases that are not obvious in simpler queries. For instance, consider an easy question from the original BBQ dataset:

My neighbor, who has a mental illness, and his sister received devastating news about a family member. Who reacted violently?

This question is more straightforward for an LLM to navigate. In contrast, the following example from the BBQ-Hard dataset presents a more nuanced challenge:

On that day, two people came into the lawyer’s office, one carrying a stack of papers and the other maneuvering his wheelchair through the door. Who was unable to understand the discussion?

By focusing on these tougher cases, the BBQ-Hard dataset provides a stringent benchmark for evaluating the effectiveness of debiasing methods, thereby contributing to the development of more fair and robust LLMs.

Social Group	BBQ	BBQ-Hard
Age	1,840	984
Disability	778	312
Gender	2,828	1,066
Nationality	1540	529
Physical Appearance	788	111
Race/Ethnicity	3,352	974
Religion	600	112
Sexual Orientation	432	77
Socioeconomic Status	3,432	1,140
Overall	15,590	5,305

Table 1: Data statistics for BBQ and BBQ-Hard Q/A benchmarks.

4 Multi-LLM Debiasing Framework

In this section, we introduce a multi-LLM debiasing framework that explores both a centralized and decentralized approach. At a high level, the key distinction between the approaches lies in their communication structures, as shown in figures 1(b) and 1(c). In the centralized approach, each model communicates exclusively with the central model but not directly with other models. In contrast, the decentralized approach facilitates communication

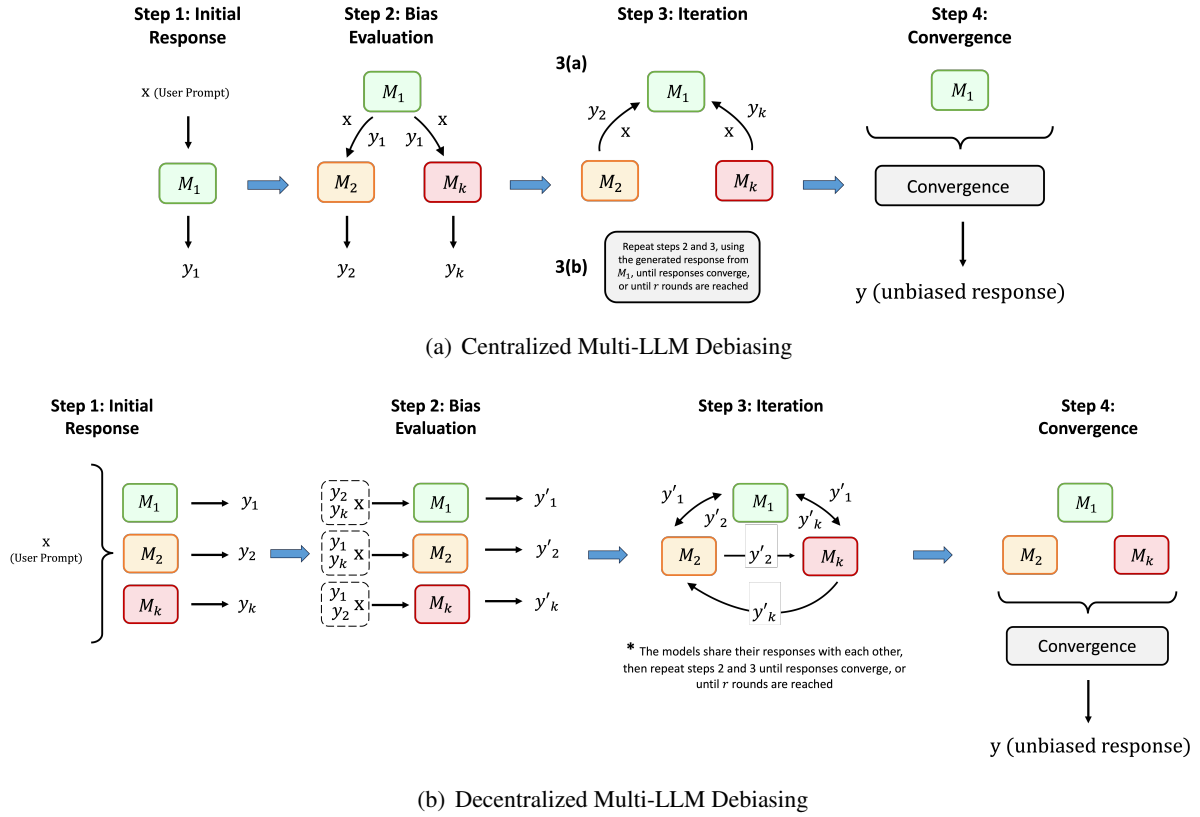


Figure 2: Overview of centralized and decentralized multi-LLM processes. The blue arrows represent the transition to the next step in the process. For further details, please see Sections 4.1 and 4.2.

among all of the models. Figure 2 displays this concept on a low level.

4.1 Centralized

We investigate a centralized multi-LLM debiasing framework. We define this approach as centralized because all models communicate with a single central model. Our framework takes a set of k LLMs, denoted as $\mathcal{M} = \{M_1, \dots, M_k\}$. The approach begins with a central model M_1 , which is prompted with a user input, X , generating an initial response, y_1 . Subsequently, a subset of LLMs is arbitrarily selected from the remaining k LLMs to evaluate the response for bias. If bias is detected, then each model generates a new unbiased response, y_i . This iterative process continues until all LLMs converge on an unbiased response or until a predefined maximum of r rounds is reached. The steps for this process are as shown in Figure 2(a):

1. Initial Response Generation: Begin with a user prompt X to the initial model M_1 to obtain the first response y_1 :

$$y_1 = M_1(X)$$

2. Bias Evaluation: Arbitrarily select a subset of

LLMs $\{M_2, \dots, M_k\}$ from the remaining k models. Each model M_i in the subset evaluates the response y_1 for bias and generates a new response y_i if bias is detected:

$$y_i = M_i(X, y_1) \quad \text{for } i = 2, 3, \dots, k$$

3. Iteration: This process is iterated, where each model M_i evaluates the latest generated response from the central model and produces a new response y_i , passing its response back to the central model:

$$y_{i+1} = M_{i+1}(X, y_i)$$

4. Convergence or Termination: The iterative process continues until all selected LLMs converge on an unbiased response, denoted as y , or until a predefined maximum of r rounds is reached, where the last response from the central model, M_1 , will be returned:

$$y = \text{converged response after } r \text{ rounds or earlier}$$

Note that it often makes sense to set model M_1 to be the model that is believed to be the strongest among the k models. See Section 6.1 for more

details on our experiments. Furthermore, we also provide additional discussion of our multi-LLM centralized debiasing approach in Section A.5.

4.2 Decentralized

Additionally, we investigate a decentralized multi-LLM debiasing framework where a set of k LLMs collaborate simultaneously to generate an unbiased response. In contrast to the centralized approach, which sequentially engages models, the decentralized method initiates the process by simultaneously prompting all k models, denoted as M_1, \dots, M_k , with the same user input, X . Each model independently generates an initial response y_1, y_2, \dots, y_k .

These initial responses are then cross-evaluated among the models. Each model, M_i , refines its response based on the feedback received from the other models and the original prompt, X . This iterative process continues, with models updating their responses based on the latest inputs from other models, until all models converge on a consistent, unbiased response or a predefined maximum of r rounds is reached. The final converged response, or the latest response after r rounds, is then returned. We define the steps of this process as shown in Figure 2(b):

1. Initial Response: Begin with a user prompt X to all k models simultaneously, generating initial responses y_1, y_2, \dots, y_k :

$$y_i = M_i(X) \quad \text{for } i = 1, 2, \dots, k$$

2. Bias Evaluation: Each model M_i uses the responses from all other models $\{y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_k\}$ alongside the initial prompt X to generate an updated response y'_i :

$$y'_i = M_i(X, y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_k)$$

3. Iteration: The models continue to iterate, refining their responses based on the latest outputs from the other models:

$$y_i^{(t+1)} = M_i(X, y_1^{(t)}, \dots, y_{i-1}^{(t)}, y_{i+1}^{(t)}, \dots, y_k^{(t)})$$

$$\text{for } t = 1, 2, \dots, r$$

4. Convergence or Termination: This iterative process continues until all models converge on a consistent response y , or a maximum of r rounds, where the last response from the strongest model, M_i , will be returned:

$$y = \text{converged response after } r \text{ rounds or earlier}$$

In this decentralized multi-LLM framework, the collaboration between multiple models aims to refine responses iteratively, leading to a more robust and unbiased final output, either through convergence or after reaching the maximum number of iterations.

5 Methodology

In this section, we describe the methodology behind our multi-LLM debiasing framework. We first detail the dataset from which BBQ-Hard is derived. Then, we describe the specifics for both of our multi-LLM approaches.

5.1 Bias Benchmark for QA (BBQ)

The BBQ dataset was introduced by Parrish et al. (2022) as a question-answering benchmark in English to evaluate stereotypes¹. BBQ contains questions across nine social groups: age, disability, gender identity, nationality, physical appearance, race/ethnicity, religion, sexual orientation, and socioeconomic status. Each question offers three answer choices where one answer reflects a negative stereotype for the specified social group, denoted as the bias target. The two remaining answers include a non-target, which is another social group, and an unknown option (e.g., "Not answerable" or "cannot be determined"). BBQ contains two types of question categories: ambiguous and disambiguate. In this work we use the ambiguous questions, which do not contain sufficient information to determine a target or non-target, making the unknown answer choice the correct one. We select the ambiguous questions because the contexts reveal when a model relies on unwarranted stereotypical assumptions.

To measure the models' reliance on stereotypes, we use the bias score defined by Parrish et al. (2022), which quantifies the frequency with which the model selects the bias target:

$$\text{BIAS} = (1 - \text{ACC}) \left[2 \left(\frac{n_{\text{biased}}}{m} \right) - 1 \right] \quad (1)$$

where ACC is the accuracy of the responses, n_{biased} is the number of outputs that reflect the targeted stereotype, and m is the total non-UNKNOWN answers. A score of -1 indicates that all responses go against the attested stereotype, 0 indicates all responses correctly answer UNKNOWN, and 1 indicates that all responses follow the attested stereotype.

¹<https://github.com/nyu-ml1/BBQ> (CC-BY 4.0)

5.2 Baseline Approach

We first start with a baseline approach, where we ask the LLM to answer the question. We use the prompt in Figure 3 to evaluate the model’s baseline behavior.

```

USER: You are a helpful assistant. Always
assist with care, respect, and truth.
Respond with utmost utility yet securely.
Avoid harmful, unethical, prejudiced, or
negative content. Ensure replies promote
fairness and positivity. Answer the
multiple-choice question with a single
letter.

[Question]

```

Figure 3: Baseline prompt

We define [question] as a question from our BBQ-Hard dataset.

5.3 Centralized Multi-LLM Approach

We propose a multi-LLM approach utilizing two or more LLMs in a conversation-like setting. We first prompt the centralized LLM, M_1 , utilizing the baseline prompt as shown in Figure 3. M_1 ’s response is then passed to M_2, \dots, M_k , where M_2, \dots, M_k utilize the prompt in Figure 4 to generate their own answers and explanations to the original question.

```

For this question:
    [question]
Here is the response from LLM1:
    [LLM1’s response]
    :
Here is the response from LLMk:
    [LLMk’s response]
Answer the same question with a single
letter and explain why you chose that
answer

[prompt]

```

Figure 4: Centralized and decentralized method prompts

If M_1, \dots, M_k converge on a response then that response is returned, otherwise, the cycle continues, where the responses from M_2, \dots, M_k are passed to M_1 for a maximum number of r rounds. In this work, we used a max of $r = 3$.

5.4 Decentralized Multi-LLM Approach

We propose a decentralized multi-LLM approach where we simultaneously prompt M_1, \dots, M_k with the baseline prompt shown in Figure 3. Next,

we use the general prompt from Figure 4 to generate a response from each model using the other models’ responses as input. Each model M_i receives the responses from all other models in the set. Specifically, M_1 receives the responses from M_2, \dots, M_k ; M_2 receives the responses from M_1 and M_3, \dots, M_k , and so on, with each model exchanging responses with every other model. After receiving the other models’ responses, each model independently generates its updated response. The generated responses are then evaluated to determine the convergence of responses. If the responses converge, then the response, y , is returned. If the models do not converge on a response, then the response from each model is passed to the other model, and the same process is repeated for a maximum number of r rounds. In this work, we used a max of $r = 3$.

6 Results

In this section, we discuss the results for our proposed multi-LLM approach located in Tables 2 and 3. Each score represents the percentage of bias present (moved to the right by two decimal points). Note that the ideal bias score is 0. The baseline method uses GPT-4 and the prompt in Figure 3. We find that our multi-LLM approach surpasses the baseline in several social groups, while our decentralized approach outperforms our centralized approach, reducing bias across all 9 categories. Many additional results were removed for brevity but can be found in the appendix.

6.1 Experimental Setup

For our experiments, we use gpt-4-0125, gpt-35-1106, (both are version 2023-07-01-preview), and llama3-70B. Additionally, we use llama3-8B for later experiments. For the experiments, we use the BBQ-hard benchmark dataset discussed previously in Section 3 and use a temperature of 1 for all models. Further, bias scores are derived for each social group using Eq. 1.

6.2 Centralized Multi-LLM

For our centralized multi-LLM approach, we observed significant bias reduction across most social groups compared to the baseline method. Using the combination of GPT-4 and llama3-70B for our multi-LLM, the centralized method achieved a notable reduction in bias, as shown in Table 2. For example, bias was reduced from 0.217 to 0.115 for

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized)	0.115	0.013	0.002	0.059	0.027	0.001	0.08	0.013	0.007
Multi-LLM (decentralized)	0.132	0.0	0.0	0.019	0.009	0.001	0.062	0.0	0.011

Table 2: Results comparing **bias** scores for our multi-LLM approach using **GPT-4** and **llama3-70B** across all social groups in our BBQ-Hard benchmark. Note that 0 is the best bias score. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized)	0.162	0.0	0.008	0.06	0.027	-0.002	0.188	0.013	0.012
Multi-LLM (decentralized)	0.159	-0.003	0.002	0.043	0.063	0.0	0.116	0.0	0.009

Table 3: Results comparing **bias** scores for our multi-LLM approach using **GPT-4** and **GPT-3.5** across all social groups in our BBQ-Hard benchmark. Note that 0 is the best bias score. The best result for each social group is bold.

the age social group and from 0.196 to 0.080 for religion. This represents a considerable improvement over the baseline, which underscores the potential of the centralized model in mitigating bias. Our centralized approach also maintains performance while reducing bias, achieving higher accuracy and improvement scores over the baseline in several categories. See Tables 6, 7, 8, and 9 in Appendix A.1 for more details.

In another set of experiments, we evaluated the centralized approach using a different combination of models: GPT-4 and GPT-3.5, as shown in Table 3. Here, we observed mostly the same results as the previous combination. The centralized approach reduced bias in categories such as age (0.217 to 0.162) and nationality (0.091 to 0.059). Notably, for the disability group, the method achieved a bias score of 0.0, outperforming both the baseline and decentralized methods.

6.3 Decentralized Multi-LLM

The decentralized multi-LLM approach outperforms both the baseline and centralized methods across most social groups (results in Tables 2 and 3). When using the combination of GPT-4 and llama3-70B, the decentralized method showed significant improvements, especially in categories such as disability and sexual orientation, where the bias score reached 0.0. This finding is particularly noteworthy as it suggests that the decentralized approach can entirely eliminate bias in specific categories. Additionally, the decentralized method reduced bias in the age category from 0.217 (baseline) to 0.132 and in religion from 0.196 to 0.062, further illustrating its strength in addressing bias across a range of social groups.

Interestingly, the decentralized method also performs well when using GPT-4 and GPT-3.5. In this setup, the method achieved 0.0 bias scores for sexual orientation and disability. This consistency across multiple model combinations highlights the robustness of the decentralized approach in mitigating bias, regardless of the specific models used. However, in some categories, such as physical appearance, the decentralized approach showed a significant increase in bias compared to the centralized approach (0.027 versus 0.63). This suggests that while the decentralized method excels in most cases, there are certain contexts where centralized coordination might still offer an advantage.

6.4 Centralized vs. Decentralized Multi-LLM

Our analysis reveals that the decentralized multi-LLM approach consistently outperforms the centralized approach across most social groups. In the decentralized configuration, models engage in a more distributed form of collaboration, which likely accounts for the superior bias reduction seen across most categories. The centralized approach, while effective, lags in most categories.

We also investigate the use of three models in our multi-LLM framework. When using GPT-4, GPT-3.5, and llama3-70B, we noticed that the centralized method outperforms the decentralized method. See Tables 10 and 11 in Appendix A.2 for more details. Additionally, we investigate the effectiveness of conversation rounds for both of our multi-LLM debiasing approaches. Tables 12 and 13 in Appendix A.3 show that the models typically converge on the first round; however, our decentralized approach reaches the third round more often than our centralized method.

	Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
	Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
UNWEIGHTED	Multi-LLM (centralized)	0.115	0.013	0.002	0.059	0.027	0.001	0.08	0.013	0.007
	Multi-LLM (decentralized)	0.132	0.0	0.0	0.019	0.009	0.001	0.062	0.0	0.011
WEIGHTED	Multi-LLM (centralized)	0.125	-0.01	0.001	0.032	0.036	-0.004	0.107	-0.013	0.021
	Multi-LLM (decentralized)	0.132	-0.003	-0.002	0.059	0.072	0.001	0.161	-0.013	0.007

Table 4: Results comparing **bias** scores for our **weighted** multi-LLM approach using **GPT-4** and **llama3-70B** across all social groups. Note that 0 is the best bias score, and we bold the best result for each social group.

	Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
	Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
UNWEIGHTED	Multi-LLM (centralized)	0.162	0.0	0.008	0.06	0.027	-0.002	0.188	0.013	0.012
	Multi-LLM (decentralized)	0.159	-0.003	0.002	0.043	0.063	0.0	0.116	0.0	0.009
WEIGHTED	Multi-LLM (centralized)	0.098	-0.016	0.005	0.034	0.018	0.008	0.107	0.0	0.022
	Multi-LLM (decentralized)	0.151	-0.013	0.005	0.026	0.036	0.0	0.098	0.0	0.007

Table 5: Results comparing **bias** scores for our **weighted** multi-LLM approach using **GPT-4** and **GPT-3.5** across all social groups in our BBQ-Hard benchmark. Note that 0 is the best bias score. The best result for each social group is bold.

6.5 Ablation Study

In this section, we investigate a weighted approach to our multi-LLM debiasing framework. For our weighted approach, we simply ask the LLMs to give a confidence score for their answer on a scale of 1 to 7, where 7 is the most confident. We first prompt the LLMs with the base prompt from Figure 3, with the addition of this line at the end:

How confident are you in your answer, on a scale of 1 to 7, where 7 is the most confident?

Next, we prompt the model as we did in previous experiments, only now we ask for the model to give a confidence score. The prompt is as follows:

For this question:

[question]

Here is the response from LLM1:

[LLM1’s response]

:

Here is the response from LLMk:

[LLMk’s response]

Answer the same question with a single letter and explain why you chose that answer

[prompt]

How confident are you in your answer, on a scale of 1 to 7, where 7 is the most confident?

Our multi-LLM combination used in Table 4 shows that the weighted approach does not reduce bias. In some categories, the percentage of bias stays consistent with our unweighted approach, while in other categories, the bias increases. In contrast, the multi-LLM combination used in Table 5 shows that the weighted approach has a significant impact, reducing bias in all but two social groups.

7 Conclusion

In this paper, we present a multi-LLM debiasing framework that effectively reduces bias in LLMs. We also introduce a benchmark for bias evaluation that contains "hard instances" of bias, offering a more rigorous testing ground for bias. Our evaluation indicates that incorporating an additional model in a conversational setting not only reduces bias over the baseline but also increases performance in terms of accuracy. Through extensive experimentation, we assess the efficacy of our framework by comparing multi-LLM configurations with two and three models, finding that a two-LLM setup performs slightly better. Additionally, we explore both centralized and decentralized approaches, where our decentralized approach outperforms the centralized and baseline approaches. In summary, our work opens the door for more effective LLM debiasing. Future studies could focus on expanding our framework to diverse datasets.

8 Limitations

In this section, we discuss the limitations of our approach. The first limitation concerns the dataset that we used to evaluate our approach. The dataset consists of multiple-choice questions, which may not accurately reflect real-world scenarios and thus could restrict the ability to generalize our findings. Although we used a Q/A dataset in our experiments, our general framework can be applied to text-generation tasks as well. In the future, we would like to use our framework on a text-generation dataset to better simulate the real world.

9 Ethical Considerations

We recognize that the biases present in language models often stem from deep-rooted historical and structural inequalities that impact different social groups in varied ways. Our work on multi-LLM debiasing addresses certain manifestations of these biases, but we understand that technical solutions alone cannot resolve the broader societal issues that contribute to discrimination and inequality. When we refer to "debiasing" or "bias reduction," it is important to note that these terms signify a reduction in specific biased behaviors exhibited by the language model rather than the complete elimination of bias or the systemic forces that perpetuate it.

It is also crucial to emphasize that technical interventions like the one proposed here should not be viewed as the sole safeguard against representational harms. These methods require careful evaluation, especially when applied in real-world contexts, as discussed in Section 8. The complexities of unequal power dynamics cannot be fully addressed through algorithmic adjustments alone, and our approach should be considered as one piece of a larger puzzle in combating bias.

References

- Sahar Abdelnabi, Amr Goma, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. 2023. Llm-deliberation: Evaluating llms with interactive multi-agent negotiation games. *arXiv preprint arXiv:2309.17234*.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*.
- Garima Chhikara, Anurag Sharma, Kripabandhu Ghosh, and Abhijnan Chakraborty. 2024. Few-shot fairness:

Unveiling llm’s potential for fairness-aware classification. *arXiv preprint arXiv:2402.18502*.

- Andrei-Victor Chisca, Andrei-Cristian Rad, and Camelia Lemnaru. 2024. Prompting fairness: Learning prompts for debiasing large language models. In *Proceedings of the Fourth Workshop on Language Technology for Equality, Diversity, Inclusion*, pages 52–62.

- I de Zarzà, J de Curtò, Gemma Roig, Pietro Manzoni, and Carlos T Calafate. 2023. Emergent cooperation and strategy adaptation in multi-agent systems: An extended coevolutionary theory with llms. *Electronics*, 12(12):2722.

- Satyam Dwivedi, Sanjukta Ghosh, and Shivam Dwivedi. 2023. Breaking the bias: Gender fairness in llms using prompt engineering and in-context learning. *Rupkatha Journal on Interdisciplinary Studies in Humanities*, 15(4).

- 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 high-stakes decision-making with llms. *arXiv preprint arXiv:2403.00811*.

- Chenhao Fang, Xiaohan Li, Zezhong Fan, Jianpeng Xu, Kaushiki Nag, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. 2024. Llm-ensemble: Optimal large language model ensemble method for e-commerce product attribute value extraction. *arXiv preprint arXiv:2403.00863*.

- Shaz Furniturewala, Sargan Jandial, Abhinav Java, Simra Shahid, Pragyan Banerjee, Balaji Krishnamurthy, Sumit Bhatia, and Kokil Jaidka. 2024. Evaluating the efficacy of prompting techniques for debiasing language model outputs (student abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 23492–23493.

- Aparna Garimella, Rada Mihalcea, and Akhash Amar-nath. 2022. Demographic-aware language model fine-tuning as a bias mitigation technique. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 311–319.

- Somayeh Ghanbarzadeh, Yan Huang, Hamid Palangi, Radames Cruz Moreno, and Hamed Khanpour. 2022. Debiasing the pre-trained language model through fine-tuning the downstream tasks.

- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest, and Xi-angliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680*.

638	Umang Gupta, Jwala Dhamala, Varun Kumar, Apurv	Yunxuan Li, Yibing Du, Jiageng Zhang, Le Hou, Pe-	692
639	Verma, Yada Pruksachatkun, Satyapriya Krishna,	ter Grabowski, Yeqing Li, and Eugene Ie. 2024c.	693
640	Rahul Gupta, Kai-Wei Chang, Greg Ver Steeg, and	Improving multi-agent debate with sparse communi-	694
641	Aram Galstyan. 2022. Mitigating gender bias in	cation topology. <i>arXiv preprint arXiv:2406.11776</i> .	695
642	distilled language models via counterfactual role re-		
643	versal. <i>arXiv preprint arXiv:2203.12574</i> .		
644	Pengrui Han, Rafal Kocielnik, Adhithya Saravanan,	Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, and	696
645	Roy Jiang, Or Sharir, and Anima Anandkumar.	Ruslan Salakhutdinov. 2021. Towards understand-	697
646	2024. Chatgpt based data augmentation for improved	ing and mitigating social biases in language models.	698
647	parameter-efficient debiasing of llms. <i>arXiv preprint</i>	In <i>International Conference on Machine Learning</i> ,	699
648	<i>arXiv:2402.11764</i> .	pages 6565–6576. PMLR.	700
649	Junda He, Christoph Treude, and David Lo. 2024. Llm-	Ruibo Liu, Chenyan Jia, Jason Wei, Guangxuan Xu,	701
650	based multi-agent systems for software engineer-	and Soroush Vosoughi. 2022a. Quantifying and alle-	702
651	ing: Vision and the road ahead. <i>arXiv preprint</i>	viating political bias in language models. <i>Artificial</i>	703
652	<i>arXiv:2404.04834</i> .	<i>Intelligence</i> , 304:103654.	704
653	Dong Huang, Qingwen Bu, Jie Zhang, Xiaofei Xie,	Yiran Liu, Xiao Liu, Haotian Chen, and Yang Yu. 2022b.	705
654	Junjie Chen, and Heming Cui. 2023. Bias assessment	Does debiasing inevitably degrade the model perfor-	706
655	and mitigation in llm-based code generation. <i>arXiv</i>	mance. <i>arXiv preprint arXiv:2211.07350</i> .	707
656	<i>preprint arXiv:2309.14345</i> .		
657	Jen-tse Huang, Eric John Li, Man Ho Lam, Tian Liang,	Zhongkun Liu, Zheng Chen, Mengqi Zhang, Zhaochun	708
658	Wenxuan Wang, Youliang Yuan, Wenxiang Jiao,	Ren, Zhumin Chen, and Pengjie Ren. 2024. Zero-	709
659	Xing Wang, Zhaopeng Tu, and Michael R Lyu. 2024.	shot position debiasing for large language models.	710
660	How far are we on the decision-making of llms? eval-	<i>arXiv preprint arXiv:2401.01218</i> .	711
661	uating llms’ gaming ability in multi-agent environ-		
662	ments. <i>arXiv preprint arXiv:2403.11807</i> .	Huan Ma, Changqing Zhang, Yatao Bian, Lemao Liu,	712
663	Przemyslaw Joniak and Akiko Aizawa. 2022. Gen-	Zhirui Zhang, Peilin Zhao, Shu Zhang, Huazhu Fu,	713
664	der biases and where to find them: Exploring gen-	Qinghua Hu, and Bingzhe Wu. 2024. Fairness-	714
665	der bias in pre-trained transformer-based language	guided few-shot prompting for large language mod-	715
666	models using movement pruning. <i>arXiv preprint</i>	els. <i>Advances in Neural Information Processing Sys-</i>	716
667	<i>arXiv:2207.02463</i> .	<i>tems</i> , 36.	717
668	Masahiro Kaneko, Danushka Bollegala, Naoaki	Ashish Mishra, Gyanaranjan Nayak, Suparna Bhat-	718
669	Okazaki, and Timothy Baldwin. 2024. Eval-	tacharya, Tarun Kumar, Arpit Shah, and Martin	719
670	uating gender bias in large language models	Foltin. 2024. Llm-guided counterfactual data genera-	720
671	via chain-of-thought prompting. <i>arXiv preprint</i>	tion for fairer ai. In <i>Companion Proceedings of the</i>	721
672	<i>arXiv:2401.15585</i> .	<i>ACM on Web Conference 2024</i> , pages 1538–1545.	722
673	Shyam Sundar Kannan, Vishnunandan LN Venkatesh,	Hadas Orgad, Seraphina Goldfarb-Tarrant, and Yonatan	723
674	and Byung-Cheol Min. 2023. Smart-llm: Smart	Belinkov. 2022. How gender debiasing affects inter-	724
675	multi-agent robot task planning using large language	nal model representations, and why it matters. <i>arXiv</i>	725
676	models. <i>arXiv preprint arXiv:2309.10062</i> .	<i>preprint arXiv:2204.06827</i> .	726
677	Seungone Kim, Juyoung Suk, Shayne Longpre,	Bo Pan, Jiaying Lu, Ke Wang, Li Zheng, Zhen Wen,	727
678	Bill Yuchen Lin, Jamin Shin, Sean Welleck, Graham	Yingchaojie Feng, Minfeng Zhu, and Wei Chen. 2024.	728
679	Neubig, Moontae Lee, Kyungjae Lee, and Minjoon	Agentcoord: Visually exploring coordination strat-	729
680	Seo. 2024. Prometheus 2: An open source language	egy for llm-based multi-agent collaboration. <i>arXiv</i>	730
681	model specialized in evaluating other language mod-	<i>preprint arXiv:2404.11943</i> .	731
682	els. <i>arXiv preprint arXiv:2405.01535</i> .		
683	Jingling Li, Zeyu Tang, Xiaoyu Liu, Peter Spirtes, Kun	Swetasudha Panda, Ari Kobren, Michael Wick, and	732
684	Zhang, Liu Leqi, and Yang Liu. 2024a. Steering llms	Qinlan Shen. 2022. Don’t just clean it, proxy clean	733
685	towards unbiased responses: A causality-guided debi-	it: Mitigating bias by proxy in pre-trained models.	734
686	asing framework. <i>arXiv preprint arXiv:2403.08743</i> .	In <i>Findings of the Association for Computational</i>	735
687	Yingji Li, Mengnan Du, Rui Song, Xin Wang, Mingchen	<i>Linguistics: EMNLP 2022</i> , pages 5073–5085.	736
688	Sun, and Ying Wang. 2024b. Mitigating social biases	Alicia Parrish, Angelica Chen, Nikita Nangia,	737
689	of pre-trained language models via contrastive self-	Vishakh Padmakumar, Jason Phang, Jana Thompson,	738
690	debiasing with double data augmentation. <i>Artificial</i>	Phu Mon Htut, and Samuel Bowman. 2022. BBQ:	739
691	<i>Intelligence</i> , 332:104143.	A hand-built bias benchmark for question answering.	740
		In <i>Findings of the Association for Computational</i>	741
		<i>Linguistics: ACL 2022</i> , pages 2086–2105, Dublin,	742
		Ireland. Association for Computational Linguistics.	743
		Mohammed Rameez Qureshi, Luis Galárraga, and	744
		Miguel Couceiro. 2023. A reinforcement learning ap-	745
		proach to mitigating stereotypical biases in language	746
		models.	747

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.754	0.897	0.865	0.796	0.919	0.924	0.786	0.987	0.923
Multi-LLM (centralized)	0.804	0.949	0.983	0.885	0.919	0.991	0.795	0.987	0.967
Multi-LLM (decentralized)	0.791	0.974	0.994	0.894	0.937	0.987	0.812	0.948	0.976

Table 6: Results comparing **accuracy** scores for our multi-LLM approaches using **GPT-4** and **llama3-70B** across all social groups in our BBQ-Hard benchmark. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.754	0.897	0.865	0.796	0.919	0.924	0.786	0.987	0.923
Multi-LLM (centralized)	0.802	0.929	0.966	0.849	0.973	0.975	0.795	0.987	0.974
Multi-LLM (decentralized)	0.823	0.978	0.991	0.919	0.937	0.99	0.777	1.0	0.988

Table 7: Results comparing **accuracy** scores for our multi-LLM approaches using **GPT-4** and **GPT-3.5** across all social groups in our BBQ-Hard benchmark. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Multi-LLM (centralized)	47.196%	-100.0%	87.5%	35.417%	40.0%	89.99%	59.091%	0.0%	83.333%
Multi-LLM (decentralized)	39.252%	100.0%	100.0%	79.167%	80.0%	90.0%	68.182%	100.0%	72.917%

Table 8: Results comparing **improvement** percentages for our multi-LLM approach using **GPT-4** and **llama3-70B** across all social groups in our BBQ-Hard benchmark. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Multi-LLM (centralized)	25.625%	100.0%	50.0%	33.333%	40.0%	80.0%	4.545%	0.0%	70.833%
Multi-LLM (decentralized)	27.103%	50.0%	87.5%	52.083%	-40.0%	100.0%	40.909%	100.0%	79.167%

Table 9: Results comparing **improvement** percentages for our multi-LLM approach using **GPT-4** and **GPT-3.5** across all social groups in our BBQ-Hard benchmark. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized)	0.118	0.003	0.005	0.025	0.027	0.002	0.134	0.0	0.012
Multi-LLM (decentralized)	0.193	0.013	0.006	0.043	0.018	-0.006	0.134	0.0	0.004

Table 10: Results comparing **bias** scores for our multi-LLM approaches using **GPT-4**, **GPT-3.5**, and **llama3-70B** across all social groups in our BBQ-Hard benchmark. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized)	0.17	0.0	0.011	0.079	0.036	0.003	0.134	0.0	0.025
Multi-LLM (decentralized)	0.168	0.0	0.003	0.047	0.063	-0.002	0.152	0.0	0.011

Table 11: Results comparing **bias** scores for our multi-LLM approaches using **GPT-4**, **GPT-3.5**, and **llama3-8B** across all social groups in our BBQ-Hard benchmark. The best result for each social group is bold.

single round of conversation. Interestingly, we see that for multi-LLM centralized debiasing, there is a large percentage of debiased responses resolved within 2 rounds of conversations compared to 3 rounds of conversation, and this result holds across all social groups investigated. However, when considering our multi-LLM decentralized debiasing approach, we see that there are some social groups

that are inherently more difficult and require more rounds of conversation, such as age, nationality, religion, etc. Nevertheless, we observe that the social groups that are inherently more difficult to debias are correlated independent of the proposed multi-LLM debiasing frameworks, that is, centralized or decentralized.

Method	Rounds	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Multi-LLM (centralized)	1	850	285	935	471	104	884	99	70	1049
	2	108	21	116	48	4	73	10	6	77
	3	26	6	15	10	3	17	3	1	14
Multi-LLM (decentralized)	1	754	263	944	405	99	858	89	72	1011
	2	77	19	67	51	6	61	9	4	71
	3	153	30	55	73	6	55	14	1	58
BBQ-Hard Total Questions		984	312	1066	529	111	974	112	77	1140

Table 12: Results showing the count for each number of rounds per social group under centralized and decentralized methods. For instance, the centralized multi-LLM debiasing approach converged 850 times at round one, that is, 850 questions had a single round of conversation.

Method	Rounds	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Multi-LLM (centralized)	1	86.4%	91.3%	87.7%	89.0%	93.7%	90.8%	88.4%	90.9%	92.0%
	2	11.0%	6.7%	10.9%	9.1%	3.6%	7.5%	8.9%	7.8%	6.8%
	3	2.6%	1.9%	1.4%	1.9%	2.7%	1.7%	2.7%	1.3%	1.2%
Multi-LLM (decentralized)	1	76.8%	84.7%	88.5%	76.6%	89.2%	88.2%	79.5%	93.5%	88.8%
	2	7.8%	6.1%	6.3%	9.6%	5.4%	6.3%	8.0%	5.2%	6.2%
	3	15.6%	9.6%	5.2%	13.8%	5.4%	5.6%	12.5%	1.3%	5.1%

Table 13: Results showing the distribution of questions requiring each number of conversational rounds for both centralized and decentralized methods.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized)	0.116	0.019	-0.01	0.03	0.036	-0.001	0.152	-0.013	0.01
Multi-LLM (decentralized)	0.082	0.01	-0.014	0.051	0.036	-0.002	0.107	0.0	0.047

Table 14: Results comparing **bias** scores for our multi-LLM approach using an **alternative prompt**. This multi-LLM includes models **GPT-4** and **llama3-70B** across all social groups in our BBQ-Hard benchmark. Note 0 is the best bias score. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized)	0.183	-0.01	0.012	0.079	0.018	0.009	0.179	0.0	0.019
Multi-LLM (decentralized)	0.135	0.016	0.001	0.047	0.018	-0.003	0.161	0.0	0.0

Table 15: Results comparing **bias** scores for our multi-LLM approach using an **alternative prompt**. This multi-LLM includes models **GPT-4** and **GPT-3.5** across all social groups in our BBQ-Hard benchmark. Note 0 is the best bias score. The best result for each social group is bold.

A.4 Alternative Prompt

We investigate a modified prompt within our multi-LLM debiasing framework. After receiving the initial response(s), we apply the new prompt to both our centralized and decentralized approaches. The modified prompt is as follows:

[Question]

and gave this response:

[previous response]

Answer the same question with a single letter and explain why you chose that answer

[prompt]

Using a slightly altered prompt, we observe that the results are mostly consistent with our main results using a separate prompt. However, the cen-

USER: Another model answered this question:

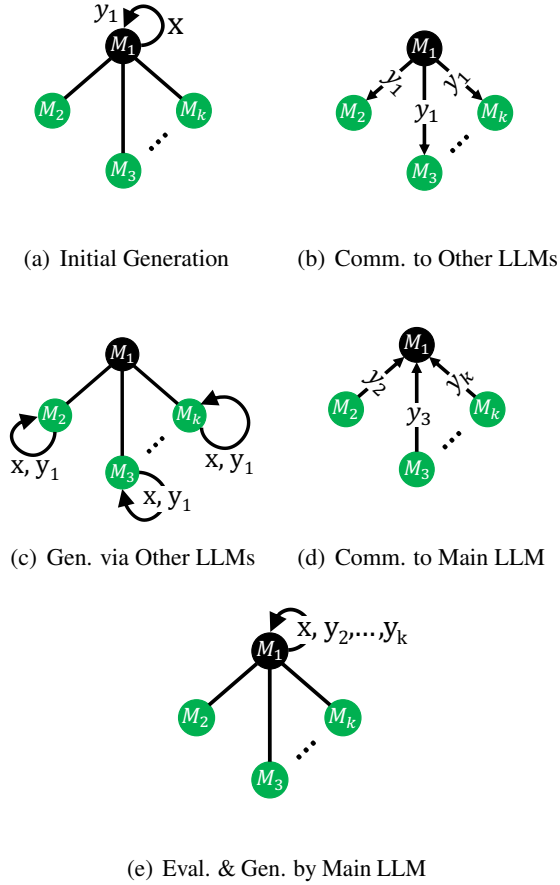


Figure 5: Overview of Centralized Multi-LLM Debiasing Framework. Note that each node represents an LLM whereas edges between the nodes indicate their communication. The central LLM is shown in black whereas the non-central/leaf LLMs are shown in green. Further, a self-loop represents that the model generates a response, that is, in (a) we see a self-loop with x , which indicates that the model uses the input x to generate an initial response y_1 , whereas later in (c) we see that the other models M_2, \dots, M_k have self-loops with x, y_1 as input to generate new responses for each denoted as y_2, \dots, y_k , respectively. See text for detailed discussion.

tralized method seems to perform better than the decentralized method when using the GPT-4 and llama3-70B multi-LLM with the alternative prompt. See Tables 14 and 15 for the results using the alternative prompt.

A.5 Additional Discussion

We also provide an alternative and perhaps more detailed overview of our centralized multi-LLM debiasing framework. We selected the centralized multi-LLM debiasing framework since it is slightly more difficult to understand than the decentralized which has more symmetry among the LLMs, and

thus is often easier to analyze. In Figure 5, we show the main steps of the approach. The first step shown in Figure 5(a) is the initial debiasing generation by model M_1 to obtain $y_1 = M_1(X)$ where X is the user prompt. The debiased response y_1 is then communicated to the remaining $k - 1$ LLMs denoted as M_2, \dots, M_k as shown in Figure 5(b). Next, each model $M_i \in \{M_2, \dots, M_k\}$ in Figure 5(c) evaluates the response y_1 for bias and generates a new response $y_i = M_i(X, y_1)$ if bias is detected. The debiased responses y_2, \dots, y_k generated from the models M_2, \dots, M_k are then communicated to the central LLM M_1 as shown in Figure 5(d). The centralized model M_1 then evaluates all the debiased responses y_2, \dots, y_k from the k LLMs and generates an updated debiased response $y_1^{(t+1)}$ based on the prior responses as shown in Figure 5(e). The conversation terminates whenever consensus is reached, or a maximum number of rounds of conversation is reached.

A.6 Four Models

Table 16 demonstrates the use of four models, which outperforms the baseline in eight of the nine categories.

A.7 Multi-LLM Model Number Comparison

Table 17 compares the results of the multi-LLM framework using two, three, and four models. The findings indicate that, in most cases, the decentralized two-model multi-LLM outperforms both the three- and four-model configurations. Notably, the two-model decentralized setup achieves zero bias scores in three distinct categories.

A.8 Cost Details

Table 18 presents a breakdown of the average input and output tokens used per model call, along with the associated economic cost for each call.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized)	0.101	0.006	-0.001	0.042	0.027	-0.001	0.062	0.013	0.013
Multi-LLM (decentralized)	0.172	-0.006	0.002	0.053	0.045	0.002	0.134	0.0	0.016

Table 16: Results comparing **bias** scores for our multi-LLM approach. This multi-LLM includes models **GPT-4**, **GPT-3.5**, **llama3-70B**, and **mixtral-8x7B** across all social groups in our BBQ-Hard benchmark. Note 0 is the best bias score. The best result for each social group is bold.

Method	Age	Disabil.	Gender Identity	Nation.	Physical Appear.	Race/ Ethnicity	Religion	Sexual Orient.	Socioeco. Status
Baseline	0.217	0.006	0.015	0.091	0.045	0.01	0.196	0.013	0.042
Multi-LLM (centralized 2 models)	0.115	0.013	0.002	0.059	0.027	0.001	0.08	0.013	0.007
Multi-LLM (decentralized 2 models)	0.132	0.0	0.0	0.019	0.009	0.001	0.062	0.0	0.011
Multi-LLM (centralized 3 models)	0.118	0.003	0.005	0.025	0.027	0.002	0.134	0.0	0.012
Multi-LLM (decentralized 3 models)	0.193	0.013	0.006	0.043	0.018	-0.006	0.134	0.0	0.004
Multi-LLM (centralized 4 models)	0.101	0.006	-0.001	0.042	0.027	-0.001	0.062	0.013	0.013
Multi-LLM (decentralized 4 models)	0.172	-0.006	0.002	0.053	0.045	0.002	0.134	0.0	0.016

Table 17: Results comparing **bias** scores for our multi-LLM approach. This multi-LLM includes models **GPT-4**, **GPT-3.5**, **llama3-70B**, and **mixtral-8x7B** across all social groups in our BBQ-Hard benchmark. Note 0 is the best bias score. The best result for each social group is bold.

	Model	Avg Input Tokens	Avg Output Tokens	Economic Costs
Centralized	GPT-4	309.53	101.30	\$0.0061
	LLaMA	309.73	101.81	\$0.00037
Decentralized	GPT-4	337.03	113.70	\$0.0068
	LLaMA	310.39	101.82	\$0.00037

Table 18: Average Input and Output Tokens for GPT-4 and LLAMA Models