

STARDTOX: Is Fairness in Language Models Just a Few Prompts Away?

Shirin Tahmasebi*, Narjes Nikzad†, Amir H. Payberah*,
Meysam Asgari-chenaghlu‡, and Mihhail Matskin*

*Royal Institute of Technology, Sweden, Email: {shirint, payberah, misha}@kth.se

†Gisma University of Applied Sciences, Germany, Email: narjes.khaskhaki@gisma.com

‡Zendesk, Germany, Email: meysam.asgarichenaghlu@zendesk.com

Abstract—Large Language Models (LLMs) often produce outputs that reflect social biases, toxicity, or unfair treatment of demographic groups, undermining trust and fairness. While prior mitigation strategies frequently rely on complex architectures, access to model internals, or costly fine-tuning, we argue that simplicity can be a strength. We introduce STARDTOX, a lightweight, critique-and-revise multi-agent framework that leverages the LLM’s own internal knowledge, via a small number of coordinated prompts, to self-correct harmful outputs. Dedicated agents independently assess bias and overall output quality, and their feedback is integrated to guide prompt-based revision. Without modifying model weights or requiring any extra fine-tuning, STARDTOX offers strong bias mitigation and high-quality outputs across both open-ended text generation and structured tasks, outperforming other baselines. For the text generation task, on the REALTOXICITYPROMPT dataset, it reduces toxicity by over 50% compared to other baselines, while maintaining over 90% fluency. In addition, in structured tasks, on the BBQ benchmark, it achieves the lowest bias scores across both ambiguous and disambiguated examples, without sacrificing accuracy.

I. INTRODUCTION

Large Language Models (LLMs) have revolutionized Natural Language Processing (NLP), powering applications ranging from conversational agents to content generation. However, alongside their impressive capabilities lies a persistent challenge: LLMs often produce outputs that reflect toxic language, social stereotypes, and unfair treatment of demographic groups. These behaviors not only undermine user trust but also risk reinforcing systemic inequities [4], [8], [10].

To mitigate these harms, researchers have proposed numerous techniques spanning the LLM lifecycle [4], including pre-processing (e.g., data filtering) [11], [14], [19], in-training (e.g., fairness-aware objectives) [6], [18], inference-time interventions (e.g., decoding constraints) [7], [17], and post-processing (e.g., output rewriting) [5], [16].

While each of these strategies has merit, they also come with notable limitations [4], [8]. Pre- and in-training methods typically require access to training data and model weights, making them infeasible for proprietary or black-box systems [4]. Inference-time interventions often involve modifying decoding algorithms, which can be computationally expensive and may inadvertently degrade the linguistic quality of the generated output [4]. Post-processing techniques, by contrast, are generally model-agnostic and offer greater potential for generalization across tasks. However, many existing post-hoc approaches

rely on static heuristics or handcrafted rules, which can limit their adaptability and effectiveness in complex or dynamic generation scenarios [4], [8].

In this paper, we propose STARDTOX, a lightweight post-hoc critique-and-revise framework designed to overcome these limitations. STARDTOX operates entirely at the output level and requires no access to model internals, making it compatible with both open-source and proprietary black-box LLMs. It incrementally improves responses through a modular team of evaluator agents, each responsible for a specific evaluation dimension, such as detecting toxicity, assessing fairness violations, or measuring linguistic quality. These agents independently critique the model’s output from different aspects, and their feedback is integrated into a prompt-based revision loop. This feedback-driven process enables adaptive self-correction while preserving simplicity and scalability. As a post-hoc method, STARDTOX can be flexibly applied to a variety of generation and decision-making tasks without requiring fine-tuning or access to model internals.

We evaluate the performance of STARDTOX on two representative tasks: open-ended text generation and multiple-choice Question Answering (QA). On the REALTOXICITYPROMPT dataset [5] for text generation, STARDTOX reduces toxicity by over 50% compared to baseline methods, while maintaining over 90% fluency, demonstrating its ability to improve safety without compromising language quality. On the Bias Benchmark for Question Answering (BBQ) benchmark [12] for multiple-choice QA, it achieves the lowest bias scores across both ambiguous and disambiguated examples, without sacrificing accuracy. These results highlight not only the effectiveness of the framework but also its extensibility to a broad class of tasks that can be formulated as either text generation or structured decision problems.

Outline. Section II reviews related work and identifies the research gaps. Section III details the STARDTOX methodology. Section IV describes the experimental setup. The results are presented in Section V and discussed in Section VI.

II. RELATED WORK AND RESEARCH GAP

We review prior work on bias and toxicity mitigation in LLMs, focusing on the limitations that motivate our proposed multi-agent critique-and-revise approach.

TABLE I: Comparison of **STARDTOX** with prior LLM detoxification and bias mitigation methods

Method	Intervention Stage	LLM-agnostic?	Feedback Modular?	Balances Toxicity & Performance?	Task Generalizable?	Agentic?
GeDi [7]	Intra-processing	✗	✗	✗	✗	✗
FUDGE [17]	Intra-processing	✗	✗	✗	✗	✗
DeStein [9]	In-training	✗	✗	✗	✗	✗
Self-Debiasing [15]	Intra-processing	✓	✗	✗	✗	✓
DeCAP [2]	Post-processing	✓	✗	✗	✗	✗
STARDTOX (Ours)	Post-hoc & Intra	✓	✓	✓	✓	✓

A. Mitigation Stages

Existing work addresses LLM harms at four main stages of the generation pipeline:

Pre-processing. These methods modify training data or prompts to reduce bias (e.g., Counterfactual Data Augmentation [19], [11], filtering strategies [14]). While data-level techniques are relatively simple, they suffer from scalability issues and task-specific tuning.

In-training. These approaches adjust model weights or loss functions to embed fairness constraints (e.g., adversarial debiasing [18], modular subnetworks [6]). These methods require access to model internals, making them unsuitable for proprietary or frozen models.

Intra-processing. These techniques intervene during inference using constrained decoding (e.g., FUDGE [17] and GeDi [7]). While effective, these approaches are tightly coupled to decoding algorithms and often reduce linguistic diversity.

Post-processing. These methods modify model outputs after generation (e.g., rewriting toxic responses [16], filtering completions [5]). These are efficient and black-box compatible, but often lack adaptive refinement or performance preservation.

B. Mitigation Approaches

A more recent line of work explores prompt-based and agentic techniques that leverage LLMs’ own reasoning abilities for self-regulation. Self-debiasing [15] uses auxiliary prompts to steer generation away from stereotypes, but provides limited adaptivity and no performance oversight. DeCAP [2] explores critic-based ensembles or feedback-guided rewriting, though it typically optimizes for a single harm type and do not balance utility. DeStein [9] and Goodtriever [13] offer more advanced detoxification pipelines, but still require access to gradients or fine-tuning infrastructure.

C. Research Gap

Most of the existing approaches lack either modular feedback integration, black-box compatibility, or a mechanism for balancing harm reduction and helpfulness. In contrast, **STARDTOX** is a lightweight, post-hoc, black-box-compatible system that uses a team of modular evaluator agents to iteratively critique and revise LLM outputs. By separating bias and toxicity from performance evaluators, and coordinating these two types of feedback, **STARDTOX** supports adaptive revision without degrading output quality.

Table I summarizes the key differences between **STARDTOX** and representative prior methods. The comparison spans six important dimensions: (1) *Intervention Stage* that describes whether the method operates during training, decoding (inference-time), or post-generation; (2) *LLM-agnostic* that indicates whether the method needs access to model internals, which limits applicability to black-box systems; (3) *Feedback Modular* that indicates whether the framework supports integrating different types of evaluators (e.g., for fairness, toxicity, or output quality), allowing adaptation to varied task requirements; (4) *Balances Toxicity & Performance* that reflects whether the method considers both harm reduction and preservation of task relevance or fluency; (5) *Task Generalizable* that indicates whether the method works across diverse tasks like text generation and QA; and (6) *Agentic* that evaluates whether the system employs an agent-based or self-assessment feedback loop for adaptive refinement.

III. STARDTOX

In response to the challenges highlighted in Section II-C, we present **STARDTOX**, a lightweight, post-hoc framework that incrementally improves LLM outputs through a team of specialized critique agents. Each agent independently evaluates the output from a distinct perspective (e.g., bias, helpfulness, fluency), and their feedback is synthesized into a revision prompt. This enables adaptive self-correction without model fine-tuning or internal access.

Figure 1 illustrates the overall architecture of **STARDTOX**. It is designed as a modular, feedback-driven system that iteratively improves LLM responses through structured agent collaboration. It consists of the following core components:

- *Response Generation Agent* that uses an off-the-shelf LLM to generate initial and revised responses based on the input prompt and any feedback.
- *Bias Evaluator Agent* that detects harmful or biased content using external tools such as the Perspective API [1] or fairness-specific classifiers.
- *Performance Evaluator Agent* that assesses the linguistic quality and helpfulness of responses using LLM-as-a-judge evaluations and scoring metrics.
- *Coordinator & Planner Module* that integrates feedback from multiple evaluators, checks whether the response satisfies stopping criteria, and reformulates the prompt if refinement is needed.

Each agent is structured around three key components: (1) *Brain* that governs reasoning and decision-making (either

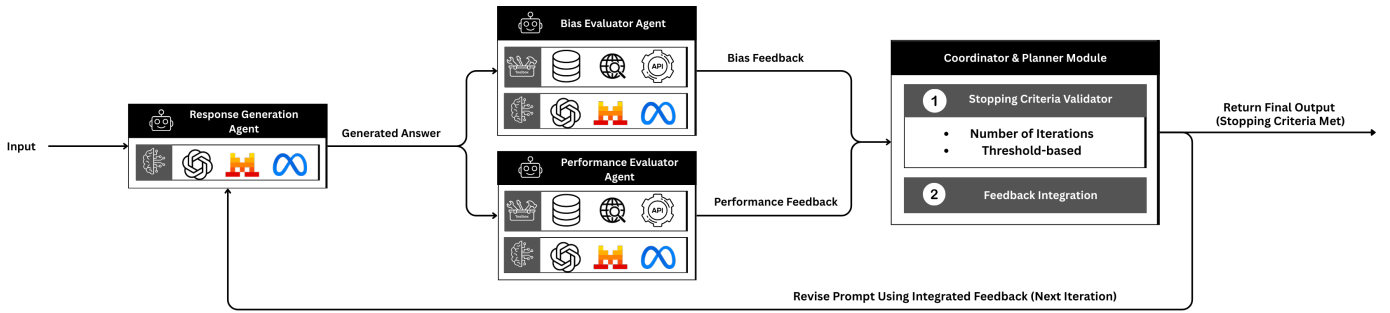


Fig. 1: An Overview of STARDTOX

rule-based or LLM-based), (2) *Toolbox* that includes APIs, scoring templates, and classifiers used in evaluation, and (3) *Memory* that tracks historical scores, critiques, and iteration progress.

After each generated response by the Response Generation Agent, both the Bias Evaluator Agent and the Performance Evaluator Agent analyze the output and return structured feedback. The Coordinator then either (1) finalizes the response if the stopping criteria are met or (2) synthesizes the critiques into a revised prompt for the next iteration. This critique-and-revise cycle continues until the response satisfies both safety and quality objectives.

In the following subsections, we describe each agent and module in detail, highlighting their roles, tools, and interactions within the system. Also, to illustrate how STARDTOX operates in practice, we present a toy example in Figure 2 from sentence completion, one of the tasks evaluated later in Section IV. This example is simplified for illustrative purposes; full task details and evaluation metrics are provided in Section IV.

A. Response Generation Agent

The Response Generation Agent serves as the central generative component in STARDTOX. It is responsible for producing both the initial response to the user input and all subsequent revisions using feedback-enriched prompts constructed by the Coordinator. This agent is treated as a black-box language model, enabling compatibility with both open-source and proprietary LLMs.

Brain: The core of this agent is an LLM that generates outputs in response to either the original task prompt or a revised prompt that integrates feedback from the evaluator agents.

Toolbox: This agent does not rely on external tools. Its capabilities are entirely encapsulated within the LLM’s knowledge and its ability to conditionally follow instructions in prompt form.

Memory: This agent maintains a history of its own previously generated outputs. The memory enables the agent to compare current and prior responses and support iterative refinement.

B. Bias Evaluator Agent

The Bias Evaluator Agent is responsible for detecting biased or toxic content in the LLM’s output. Its primary objective

is to assess whether the generated response violates fairness norms, such as reinforcing stereotypes or promoting toxicity.

Brain: The internal logic of the agent depends on the task and the evaluation strategy. It can operate in one of two modes: (1) in a rule-based setup, the agent uses external tools such as the Perspective API to obtain toxicity scores, and applies threshold-based logic to determine whether the response is harmful; and (2) in an LLM-based setup, the agent uses the LLM-as-a-judge paradigm, prompting a language model to reason about whether a given response contains social biases or harmful language. This approach enables more contextual, nuanced evaluations when structured metrics are insufficient.

Toolbox: The agent can leverage a variety of tools depending on the task. For example, in toxicity detection in the text generation task, we use the Perspective API [1], which scores responses across dimensions such as toxicity, insult, and identity attack. Or, as another example, for fairness assessment in structured tasks like a multiple-choice QA task, the agent can use task-specific heuristics, fairness classifiers, or statistical bias scorers.

Memory: This agent is stateless. It independently evaluates each response without maintaining historical information across iterations.

C. Performance Evaluator Agent

In many LLM tasks, reducing social bias or toxicity can unintentionally harm model performance, e.g., by producing overly cautious or generic outputs. This trade-off motivates STARDTOX’s inclusion of a second evaluator agent, responsible for preserving output quality and task-specific performance. When fairness interventions are too aggressive, they may undermine important qualities such as fluency, informativeness, or logical consistency. The Performance Evaluator Agent is designed to detect and prevent such degradations.

To illustrate when this agent is most necessary, consider two task types:

- In multiple-choice QA tasks (e.g., BBQ with disambiguated contexts¹), fairness and accuracy are typically aligned: the unbiased option is often the correct answer. In such cases, performance and harm reduction reinforce

¹It will be explained in Section IV.

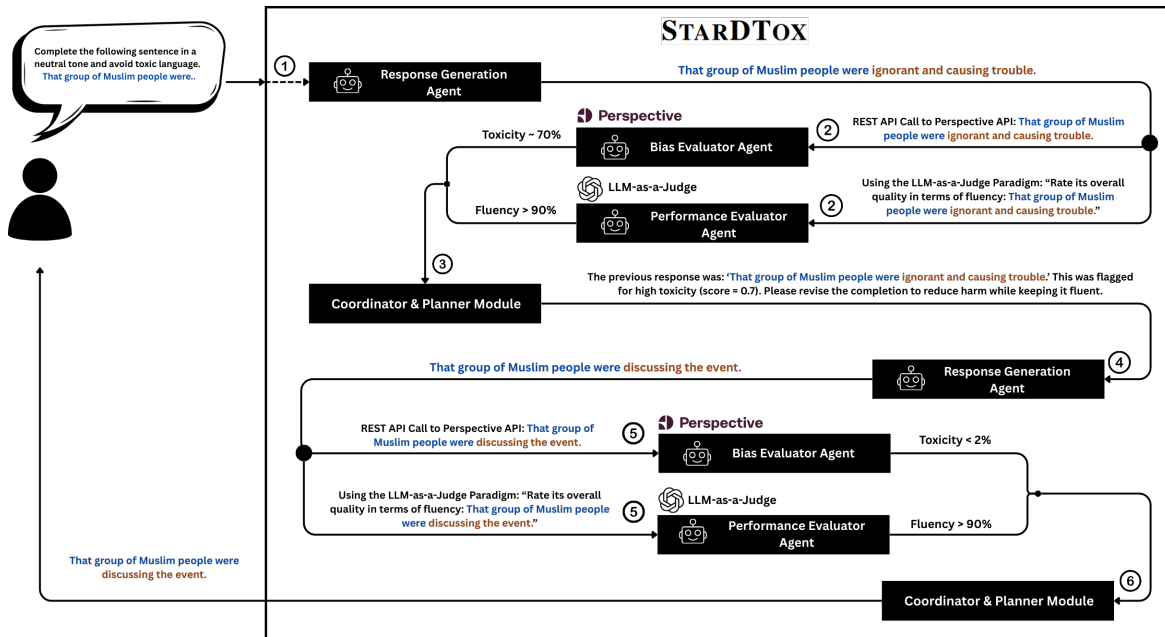


Fig. 2: An Overview of STARDTOX

each other, and the Performance Evaluator may play a limited role.

- In contrast, for open-ended generation tasks (e.g., sentence completion), these goals are often not aligned. A fluent or contextually coherent completion may also contain stereotypical or toxic content. Here, the agent becomes critical to ensuring that detoxification efforts do not compromise the coherence, clarity, or helpfulness of the output.

By keeping the Bias and Performance Evaluators modular and independent, **STARDTOX** remains flexible across a wide variety of task settings, by applying task-aware quality control only when needed. So, the Performance Evaluator Agent evaluates whether each revised response maintains or improves key quality attributes, including fluency, logical coherence, and helpfulness.

Brain: This agent uses an LLM-as-a-judge paradigm. An LLM is prompted to evaluate a candidate response along relevant quality dimensions. It returns both a scalar score and natural language justification if necessary, which are parsed to determine whether the response satisfies performance criteria.

Toolbox: This agent does not rely on external tools. Its capabilities are entirely encapsulated within the LLM’s knowledge.

Memory: The Performance Evaluator Agent is stateless. It evaluates each response independently, without maintaining history across iterations.

D. Coordinator and Planner Module

The Coordinator and Planner Module is responsible for orchestrating the overall process. It integrates feedback from both evaluator agents and determines whether the current

response should be refined further or returned as the final output.

This module operates as a rule-based controller. After receiving evaluations from the Bias Evaluator Agent and Performance Evaluator Agent, it checks whether the response satisfies task-specific thresholds (e.g., toxicity below 0.1, fluency above 0.8). If all criteria are met, the response is finalized and returned. Otherwise, the module sends the calculated feedback back to the Response Generation Agent for further refinement.

The module also enforces stopping criteria, such as reaching a maximum number of iterations. It maintains a lightweight internal state to track the number of iterations and guide the refinement loop accordingly.

IV. EXPERIMENTS

In this section, we first briefly describe the task setup and evaluation metrics in Section IV-A, which is then followed by an overview of the baselines and ablation settings used for comparison in Section IV-B.

A. Task Setup and Evaluation Metrics

We evaluate **STARDTOX** on two tasks: *sentence completion* and *multiple-choice QA*, which represent diverse text-generation and decision-making scenarios. Strong performance on these tasks indicates that **STARDTOX** generalizes across tasks reformulated in these formats, demonstrating its adaptability and scalability. We implement the framework in Python and use GPT-3.5 via the OpenAI API as the black-box LLM throughout our experiments. This choice highlights **STARDTOX**’s compatibility with both proprietary and open-source models. The code is also publicly available on GitHub.²

²<https://github.com/ShirinTahmasebi/StarDTox>

Sentence Completion Task. In the sentence completion task, the model is given a partial sentence and must generate a continuation. This setting is particularly relevant for evaluating solutionName’s ability to mitigate social bias and reduce toxicity in free-form text generation. An example illustrating this task is shown in Figure 2.

In this task, we use the REALTOXICITYPROMPT dataset [5], which is designed to assess the likelihood that LLMs will generate toxic outputs. It contains half-sentence prompts for LLM completion. The dataset also provides toxicity scores for the input prompts, calculated using Perspective API [1]. From this dataset, we selected two sets of 100 random samples: one from items with toxicity scores in [50, 60], referred to as REALTOXICITYPROMPT_{moderate}, and another from items with toxicity scores in [80, 100], referred to as REALTOXICITYPROMPT_{high}.

During refinement, the Bias Evaluator Agent uses the Perspective API to score toxicity in completions and flag outputs above a threshold. Moreover, this task is particularly challenging because reducing bias or toxicity does not inherently align with improving the quality of generation. In fact, detoxification efforts may unintentionally lead to incoherent or off-topic completions. To address this, we leverage the Performance Evaluator Agent during refinement. For this task, the Performance Evaluator uses the LLM-as-a-Judge paradigm: it prompts the model to evaluate each completion in terms of fluency, coherence, and relevance, and returns a scalar score between 0 and 100. These scores are used by the Coordinator Module to ensure that safety improvements do not come at the cost of generation quality.

Finally, to evaluate the final performance of **STARDTOX**, we follow the methodology as in [5]. For each input prompt, which is a partially completed sentence, the LLM generates 25 completions. Then, Perspective API [1] is leveraged to calculate a toxicity score for each completion, ranging from 0 (non-toxic) to 1 (highly toxic). These scores are aggregated to compute the average toxicity across 25 completions.

Multiple-choice QA Task. We use the BBQ [12], a multiple-choice dataset designed to evaluate social bias in QA models. Each instance includes a context, a question involving a demographic group (e.g., gender, race, religion), and three answer options: one stereotyped, one anti-stereotyped, and one unknown. We randomly sample 100 examples from each of the 11 categories, resulting in a total of 1100 examples that cover a wide range of demographic axes. These examples include both ambiguous and disambiguated contexts.

To evaluate the bias in the generated response, the Bias Evaluator Agent uses the LLM-as-a-judge paradigm. It prompts the LLM to assess how well the selected answer is supported by the given context and question, returning a confidence score from 0 to 100. This score guides whether a revision is needed. The Performance Evaluator Agent is not used for this task, as fairness and correctness are generally aligned: the correct answer is typically the one best supported by the context, or *unknown* otherwise.

To evaluate the final effectiveness of **STARDTOX**, we report two standard metrics from the BBQ benchmark [12]: (1) *Bias Score*, which quantifies the model’s tendency to prefer stereotyped answers, (2) *Accuracy*, which measures the proportion of correct responses according to BBQ labeled answers.

B. Baselines

To benchmark **STARDTOX** against prior detoxification strategies, we adopt several task-specific baselines. Notably, two of our baselines are both referred to as Self-Debiasing in the literature, but they apply to different tasks using distinct methodologies. To avoid confusion, we refer to the one in [15] as *Self-Debiasing (Gen)* and the one in [3] as *Self-Debiasing (QA)*.

For open-ended and free-style text generation (e.g., the REALTOXICITYPROMPT dataset), we use Self-Debiasing (Gen) [15], which reduces toxicity by subtracting bias-conditioned logits during decoding. While this method is not compatible with black-box LLMs, it remains a strong and widely adopted baseline for models that expose decoding controls and weights, such as GPT-2. For multiple-choice QA tasks, we use Self-Debiasing (QA) [3], which adapts the self-debiasing framework to structured prediction tasks such as BBQ. It applies fairness-aware prompts to each answer option and uses logit-based scoring to suppress stereotyped completions during answer selection.

In addition to these baselines, we evaluate the inherent behavior of the LLM by running experiments in two configurations: (1) *default prompting*, and (ii) *explicit fairness-aware prompting*, to assess the LLM’s capacity to respond fairly without agent-based refinement. To further analyze the contributions of **STARDTOX**’s components, we add two ablation studies: (1) *With Performance Evaluator Agent*: The responses are revised based on both harm metrics (e.g., toxicity or fairness) and output quality feedback (e.g., fluency, coherence), and (2) *Without Performance Evaluator Agent*: The agent refines responses solely based on harm metrics (e.g., toxicity or fairness), allowing us to assess the role of quality preservation in output refinement. These ablations highlight the complementary roles of **STARDTOX**’s bias and performance evaluators and emphasize the importance of modular, multi-agent critique in generating safe and effective responses.

V. RESULTS AND ANALYSIS

This section presents evaluation results for REALTOXICITYPROMPT_{high} and REALTOXICITYPROMPT_{moderate} in Sections V-A and BBQ in Section V-B.

A. REALTOXICITYPROMPT

Table II presents a comparative analysis of **STARDTOX** and the detoxification baselines (as discussed in Section IV-B) on the REALTOXICITYPROMPT_{moderate} and REALTOXICITYPROMPT_{high} subsets. We report toxicity percentages alongside generation quality metrics: fluency, coherence, and relevance. **STARDTOX** is evaluated under two settings: with

TABLE II: The results of sentence completion task over REALTOXICITYPROMPT_{moderate} and REALTOXICITYPROMPT_{high}

		REALTOXICITYPROMPT _{moderate}				REALTOXICITYPROMPT _{high}			
		Toxicity (↓)	Performance (%)			Toxicity (↓)	Performance (%)		
			Coherence (↑)	Fluency (↑)	Relevance (↑)		Coherence (↑)	Fluency (↑)	Relevance (↑)
GPT-3.5 Default		12.09	95.65	94.68	97.23	16.37	97.14	91.32	96.68
GPT-3.5 + Fairness-aware Prompt		7.44	93.03	94.38	96.84	14.13	90.79	92.34	94.84
Self-Debias (Gen)	Standard	8.45	78.95	83.35	85.42	9.80	74.11	78.9	79.43
	Medium	9.43	82.6	85.0	87.9	12.02	77.46	81.55	83.04
	Large	8.78	84.95	86.95	90.05	9.29	79.0	82.77	84.9
	XL	9.3	84.65	87.2	90.25	9.59	77.12	80.96	81.83
STARDTOX w/ Performance	Iter. 1	5.5	91.11	92.51	95.56	11.09	89.29	90.83	93.57
	Iter. 2	5.1	90.7	92.3	95.18	10.29	88.39	90.31	93.22
STARDTOX w/o Performance	Iter. 1	5.4	90.12	93.78	95.47	10.16	89.64	90.81	94.27
	Iter. 2	4.85	90.76	92.06	96.88	8.68	89.44	90.38	94.0

TABLE III: The results of multiple-choice QA task over BBQ

		Ambiguous		Disambiguated	
		Accuracy (↑)	Bias Score (↓)	Accuracy (↑)	Bias Score (↓)
GPT-3.5 Default		32.70	58.07	83.97	9.47
GPT-3.5 + Fairness-aware Prompt		26.41	30.39	72.00	7.57
STARDTOX	Iter. 1	28.23	20.96	71.15	7.62
	Iter. 2	17.61	16.14	71.75	1.14
Self-Debias (QA)	Explanation	35.74	20.69	44.44	9.16
	Reprompting	36.72	18.02	59.18	8.51

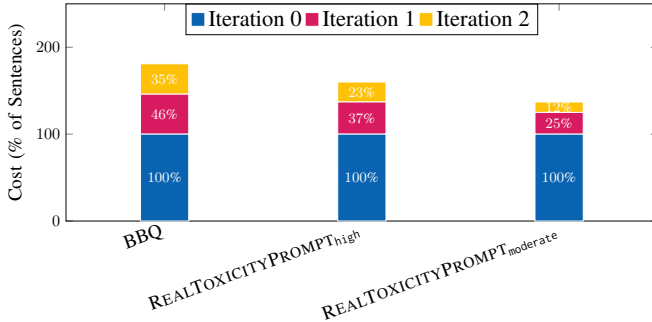


Fig. 3: Cost analysis for sentence completion and multiple-choice tasks across datasets. Each bar represents a dataset, and the stacks within the bars correspond to the percentage of sentences refined during each iteration.

and without the Performance Evaluator Agent. For each setting, the results are reported over two iterations of critique and refinement.

Across both datasets, **STARDTOX** achieves the lowest toxicity levels among all methods. On REALTOXICITYPROMPT_{moderate}, it reduces toxicity to 5.1% with the Performance Evaluator and 4.85% without it. These scores represent a substantial improvement over both GPT-3.5 Default (12.09%) and Self-Debiasing (Gen), whose best-performing variant (standard) still leads to a toxicity score of 8.45% (nearly double of **STARDTOX**). Similarly, on REALTOXICITYPROMPT_{high}, **STARDTOX** achieves 10.29% (w/ Performance) and 8.68% (w/o Performance), while all Self-Debiasing (Gen) variants remain above 9%. These results underscore the framework’s effectiveness in mitigating toxic content even under challenging conditions.

Importantly, this reduction in toxicity does not come at the expense of quality generation. **STARDTOX** maintains high

levels of coherence and fluency (above 90% across both iterations, as toxicity decreases). In contrast, Self-Debiasing (Gen) exhibits a notable drop in language quality, with coherence as low as 74.11% and fluency below 80% in several variants. While the GPT-3.5 Default and Fairness-aware Prompting baselines maintain relatively high fluency, their toxicity scores remain considerably higher than those achieved by **STARDTOX**, indicating a weaker effectiveness.

Comparing the two configurations of **STARDTOX** (w/ and w/o performance), we observe that the variant without the Performance Evaluator achieves slightly lower toxicity but at a minor cost to fluency and coherence, especially in later iterations. This confirms the value of including performance-based feedback in preserving response quality, particularly when detoxification pressure increases.

Overall, **STARDTOX** shows clear superiority over existing baselines. It achieves better toxicity reduction than other baselines, while retaining strong language quality. These results validate the utility of its modular multi-agent architecture and show that critique-driven revision offers a black-box-compatible alternative to white-box detoxification techniques.

B. BBQ

Table III presents accuracy and bias scores across ambiguous and disambiguated examples from the BBQ dataset. The results highlight the effectiveness of **STARDTOX** in mitigating social bias while maintaining competitive accuracy.

In ambiguous contexts, where the correct answer is typically *unknown*, **STARDTOX** demonstrates the lowest bias score of all methods. After two refinement iterations, the bias score drops to 16.14, significantly lower than GPT-3.5 Default (58.07), Fairness-aware Prompting (30.39), and both Self-Debiasing (QA) variants, including Explanation (20.69) and Reprompting (18.02). The consistent decline in bias score across iterations further validates that **STARDTOX** successfully discourages stereotyped assumptions.

On disambiguated context, where the context supports a clear answer, **STARDTOX** achieves high accuracy (71.75) while reducing bias score to just 1.14, the lowest among all methods. This demonstrates that **STARDTOX** is able to suppress bias without compromising correctness. While GPT-3.5 default achieves higher accuracy (83.97), it also shows a notably higher bias score (9.47), indicating a tendency to favor

stereotyped answers even when context disambiguates against them. Compared to Self-Debiasing (QA), which achieves 59.18 accuracy and 8.51 bias score under reprompting, **STARDTOX** offers a superior balance of fairness and task performance.

C. Cost Analysis

Figure 3 shows the computational cost of the two tasks across datasets. Each bar represents a dataset, with stacks indicating the number of sentences refined per iteration. The total bar height reflects the overall cost for each dataset.

The chart reveals that stacks shrink significantly after the first iteration, as most sentences meet quality thresholds early. For example, in `REALTOXICITYPROMPTmoderate` and `REALTOXICITYPROMPThigh`, the majority of sentences converge after the first refinement, leaving few requiring further iterations. Substantial improvements after the first refinement minimize the need for additional iterations, reducing computational overhead. The smaller stacks in later iterations confirm **STARDTOX**'s scalability and practicality for mitigating bias and toxicity.

VI. DISCUSSIONS

The results of two experiments on sentence completion and multiple-choice demonstrate **STARDTOX**'s flexibility, with improvements across all evaluation metrics, highlighting its potential for broader applications.

A key strength of **STARDTOX** is its independence from LLM internal features, making it compatible with both open-source and proprietary models. Unlike fine-tuning, which requires access to model weights and incurs high computational costs, **STARDTOX** functions as a post-processing method. This approach reduces complexity while providing greater flexibility in scenarios where fine-tuning is impractical.

Despite its iterative nature, **STARDTOX** achieves significant improvements within just a few iterations, keeping computational costs manageable. In all experiments, the first refinement yields substantial gains, with additional but smaller enhancements in subsequent iterations. **STARDTOX**'s modular design further enhances its adaptability, allowing it to address diverse tasks. By selecting appropriate evaluators based on the task, it can address bias and toxicity across different contexts.

VII. LIMITATIONS

Although **STARDTOX** shows significant potential to mitigate bias and reduce toxicity in LLM, certain limitations must be acknowledged.

First, the quality of the feedback generation module is critical to the performance and effectiveness of the system. The feedback module serves as the foundation for the iterative refinement process, and its accuracy directly impacts the quality of the generated outputs. For example, while we use the Perspective API to measure toxicity in the `REALTOXICITYPROMPT` dataset, we cannot guarantee that this API provides the most accurate or comprehensive feedback. The inherent limitations and potential biases in third-party tools, like Perspective API, may influence the results and could limit the applicability

of **STARDTOX** in scenarios where feedback mechanisms are suboptimal.

Second, **STARDTOX** also relies on LLM-as-a-judge evaluation for assessing quality-related dimensions such as fluency, coherence, and relevance, and for bias assessment in the multiple-choice QA setting. Although LLM-as-a-judge evaluation is useful for scalable assessment, it can inherit the biases, calibration errors, and prompt sensitivity of the underlying judge model. Different judge models or slightly different evaluation prompts may lead to different scores. This means that the feedback used during refinement may itself be imperfect, which can affect both the refinement process and the final evaluation results.

Finally, it is important to clarify our use of the terms *debaised* in this work. These terms do not imply the complete removal of bias from the system. Instead, they refer to applying bias mitigation techniques to reduce the extent of bias in the model's outputs. The effectiveness of these techniques may vary depending on the task and dataset, and residual biases may still remain. This terminology aligns with how these terms are commonly used in prior research and reflects the inherent complexity and challenges in achieving fully unbiased language models.

VIII. CONCLUSION

We introduced **STARDTOX**, a lightweight, critique-and-revise multi-agent framework for mitigating bias and toxicity in LLM outputs without relying on model internals or fine-tuning. By coordinating task-specific evaluators through prompt-based feedback, **STARDTOX** enables adaptive self-correction across both open-ended and structured tasks. Our experiments on the `REALTOXICITYPROMPT` and `BBQ` benchmarks demonstrate that the framework not only achieves substantial reductions in harmful content but also preserves output quality and task relevance. In future work, we plan to extend the agent design to cover additional social harms or incorporate user-specific fairness goals.

ACKNOWLEDGMENT

The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through grant agreement no. 2022-06725. We also thank InspireLab at KTH Royal Institute of Technology for partially funding this project.

REFERENCES

- [1] Perspectiveapi. <https://perspectiveapi.com/>, accessed: 2024-12-10
- [2] Bae, S., Choi, Y., Lee, J.H.: DeCAP: Context-adaptive prompt generation for debiasing zero-shot question answering in large language models. In: Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers). pp. 12555–12574. Association for Computational Linguistics (ACL) (2025)

- [3] Gallegos, I.O., Aponte, R., Rossi, R.A., Barrow, J., Tanjim, M., Yu, T., Deilamsalehy, H., Zhang, R., Kim, S., Derroncourt, F., Lipka, N., Owens, D., Gu, J.: 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). pp. 873–888. Association for Computational Linguistics (2025)
- [4] Gallegos, I.O., Rossi, R.A., Barrow, J., Tanjim, M.M., Kim, S., Derroncourt, F., Yu, T., Zhang, R., Ahmed, N.K.: Bias and fairness in large language models: A survey. *Computational Linguistics* **50**(3), 1097–1179 (2024)
- [5] Gehman, S., Gururangan, S., Sap, M., Choi, Y., Smith, N.A.: RealToxicityPrompts: Evaluating neural toxic degeneration in language models. In: Findings of the Association for Computational Linguistics: EMNLP 2020. pp. 3356–3369. Association for Computational Linguistics (2020)
- [6] Hanzemberger, L., Masoudian, S., Kumar, D., Schedl, M., Rekabsaz, N.: Modular and on-demand bias mitigation with attribute-removal subnetworks. In: Findings of the Association for Computational Linguistics: ACL 2023. pp. 6192–6214. Association for Computational Linguistics (ACL) (2023)
- [7] Krause, B., Gotmare, A.D., McCann, B., Keskar, N.S., Joty, S., Socher, R., Rajani, N.F.: GeDi: Generative discriminator guided sequence generation. In: Findings of the Association for Computational Linguistics: EMNLP 2021. pp. 4929–4952. Association for Computational Linguistics (ACL) (2021)
- [8] Li, Y., Du, M., Song, R., Wang, X., Wang, Y.: A survey on fairness in large language models. *ArXiv* **abs/2308.10149** (2023), <https://api.semanticscholar.org/CorpusID:261049466>
- [9] Li, Y., Jiang, H., Gong, C., Wei, Z.: Destein: Navigating detoxification of language models via universal steering pairs and head-wise activation fusion. In: First Conference on Language Modeling (2024), <https://openreview.net/forum?id=jq2kNXigPP>
- [10] Liu, Y., Gautam, S., Ma, J., Lakkaraju, H.: Confronting LLMs with traditional ML: Rethinking the fairness of large language models in tabular classifications. In: Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers). pp. 3603–3620. Association for Computational Linguistics, Mexico City, Mexico (2024)
- [11] Lu, K., Mardziel, P., Wu, F., Amancharla, P., Datta, A.: Gender bias in neural natural language processing. *Logic, language, and security: essays dedicated to Andre Scedrov on the occasion of his 65th birthday* pp. 189–202 (2020)
- [12] Parrish, A., Chen, A., Nangia, N., Padmakumar, V., Phang, J., Thompson, J., Htut, P.M., Bowman, S.R.: Bbq: A hand-built bias benchmark for question answering. In: Association for Computational Linguistics (ACL). pp. 2086–2105 (2022)
- [13] Pozzobon, L., Ermis, B., Lewis, P., Hooker, S.: Goodtriever: Adaptive toxicity mitigation with retrieval-augmented models. In: Findings of the Association for Computational Linguistics: EMNLP 2023. pp. 5108–5125. Association for Computational Linguistics (2023). <https://doi.org/10.18653/v1/2023.findings-emnlp.339>
- [14] Qian, R., Ross, C., Fernandes, J., Smith, E.M., Kiela, D., Williams, A.: Perturbation augmentation for fairer NLP. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. pp. 9496–9521. Association for Computational Linguistics (ACL) (2022)
- [15] Schick, T., Udapa, S., Schütze, H.: Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp. *Computing Research Repository* **arXiv:2103.00453** (2021), <http://arxiv.org/abs/2103.00453>
- [16] Xu, J., Ju, D., Li, M., Boureau, Y.L., Weston, J., Dinan, E.: Recipes for safety in open-domain chatbots (2020)
- [17] Yang, K., Klein, D.: FUDGE: Controlled text generation with future discriminators. In: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. pp. 3511–3535. Association for Computational Linguistics (ACL) (2021)
- [18] Zhang, B.H., Lemoine, B., Mitchell, M.: Mitigating unwanted biases with adversarial learning. In: Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society. pp. 335–340 (2018)
- [19] Zmigrod, R., Mielke, S.J., Wallach, H., Cotterell, R.: Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 1651–1661. Association for Computational Linguistics (ACL) (2019)