

Context-Aware Counterfactual Data Augmentation for Gender Bias Mitigation in Language Models

Shweta Parihar
University of Illinois at Chicago
spari@uic.edu

Natalie Parde
University of Illinois at Chicago
parde@uic.edu

Guangliang Liu
Michigan State University
liuguan5@msu.edu

Lu Cheng
University of Illinois at Chicago
lucheng@uic.edu

ABSTRACT

A challenge in mitigating social bias in fine-tuned language models (LMs) is the potential reduction in language modeling capability, which can harm downstream performance. Counterfactual data augmentation (CDA), a widely used method for fine-tuning, highlights this issue by generating synthetic data that may align poorly with real-world distributions or creating overly simplistic counterfactuals that ignore the social context of altered sensitive attributes (e.g., gender) in the pretraining corpus. To address these limitations, we propose a simple yet effective context-augmented CDA method, *Context-CDA*, which uses large LMs to enhance the diversity and contextual relevance of the debiasing corpus. By minimizing discrepancies between the debiasing corpus and pretraining data through augmented context, this approach ensures better alignment, enhancing language modeling capability. We then employ uncertainty-based filtering to exclude generated counterfactuals considered low-quality by the target smaller LMs (i.e., LMs to be debiased), further improving the fine-tuning corpus quality. Experimental results on gender bias benchmarks demonstrate that *Context-CDA* effectively mitigates bias without sacrificing language modeling performance while offering insights into social biases by analyzing distribution shifts in next-token generation probabilities.

1. INTRODUCTION

Language models (LMs) have achieved remarkable success in generating human-like text across a wide range of applications, from chatbots [22] to translation [51]. However, these models often inherit and amplify biases present in their training data, which can lead to outputs that reinforce stereotypes or perpetuate harmful prejudices. These biases partly stem from the massive datasets used to train LMs, which frequently reflect societal imbalances, discriminatory language, and historical injustices.

Despite promising results in mitigating bias in LMs, a fundamental drawback of current debiasing methods is their potential to harm the core modeling abilities of LMs [8]. These methods often rely on strategies such as removing or altering biased data [49] or introducing controlled outputs [32], which can reduce the model’s exposure to natural linguistic patterns. By filtering or distorting the underlying data, these techniques can hinder the model’s ability to grasp the subtleties of language, leading to less fluent and contextually inaccurate text generation. This trade-off underscores the challenge of reducing bias while maintaining language proficiency.

Take Counterfactual Data Augmentation (CDA) [19; 52] for gender bias as an example. CDA works by altering gender attributes to generate counterfactual examples, which are then used during fine-tuning to reduce bias in model predictions. While CDA is effective in reducing bias, it often leads to a degradation in LMs’ language modeling ability [8; 7; 29]. A major reason for this is the synthetic data generated by CDA may not align well with real-world data distributions and can create overly simplistic counterfactuals that fail to account for the social context of altered sensitive attributes (e.g., gender) embedded in the pre-training corpus.

To overcome these challenges, we propose a simple yet effective approach, *Context-CDA*, that leverages the generative capabilities of large LMs to augment the context used for debiasing. Particularly, larger LMs trained on vast and diverse corpora can generate context-rich counterfactual examples that are not only more diverse but also closely resemble natural language patterns (see an example in Figure 1). This context-aware data augmentation process helps avoid the severe distortions often introduced by traditional CDA techniques, where overly simplified examples fail to consider the social context of altered sensitive attributes. Our approach thus enables a more comprehensive debiasing process, enhancing fairness while also preserving the language modeling abilities of LMs.

However, the smaller LMs to be debiased often struggle to learn from text generated by larger LMs. Samples deemed challenging for the target LMs may hurt their language modeling capabilities [12]. To improve the performance of the model and reduce the impact of low-quality generations, we further propose using an uncertainty-based filtering strategy. Particularly, we leverage semantic entropy [15] to quantify the uncertainty of each augmented sample in our corpus. Semantic entropy advances other uncertainty estimation approaches (e.g., likelihoods) as it incorporates linguistic invariances created by shared meanings. High semantic entropy indicates that the target LM finds the sentence uncertain, ambiguous, or difficult to interpret due to multiple possible meanings or unpredictable word choices. We systematically filter out the counterfactual generations with the largest semantic entropy. Therefore, during fairness fine-tuning, the target LMs are trained on text that is clearer and less prone to misinterpretation. This filtering also ensures that the model can focus on learning meaningful and unambiguous patterns from the data, enhancing its ability to generalize and perform well in real-world applications.

Our contributions are as follows: **Method:** We develop *Context-CDA*, a new approach that produces context-rich, gender-flipped sentences. We enhance data quality for target LMs by using semantic entropy filtering to exclude ambiguous and irrelevant coun-

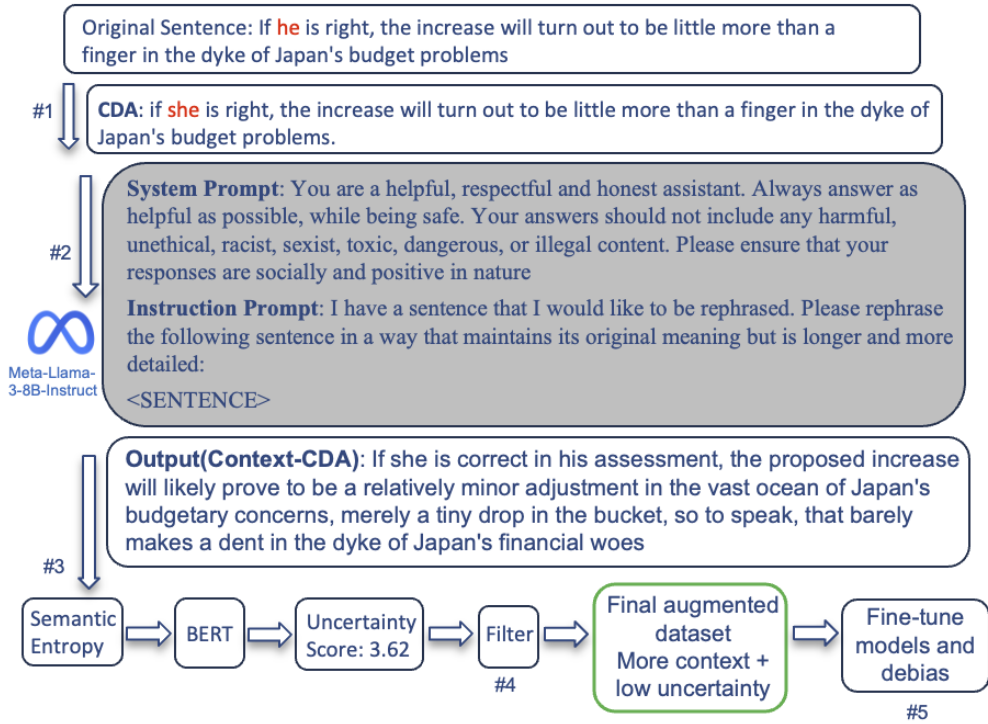


Figure 1: An illustration of the proposed *Context-CDA* pipeline using the StereoSet benchmark [23] with BERT [5]. Step 1: Flip the gender words. Step 2: Use a larger LM (e.g., Llama-3-8B-Instruct [9]) with the system and instruction prompt to get the augmented data. Step 3: Use the target small LM (e.g., BERT) to calculate the semantic entropy of augmented data. Step 4: Filter the counterfactuals based on the semantic entropy. Step 5: Debias the target small LM.

terfactuals for fine-tuning. **Experiments:** Our extensive evaluations show that *Context-CDA* preserves language modeling ability while reducing bias more effectively than traditional CDA methods (which may compromise on fluency). Semantic entropy-based filtering of low-quality counterfactuals further enhances debiasing effectiveness and maintains language fluency. Analysis of next-token distribution indicates a better gender balance and token diversity, leading to reduced skew towards male tokens, and increasing robustness. Comprehensive evaluation across five diverse model architectures, including BERT and DistilBERT (encoder-only models), T5 (encoder-decoder generative model), GPT-2, and Llama-3-1B (causal decoder-only generative model), demonstrates that these findings are consistent and robust across both discriminative and generative systems, validating the true model-agnostic nature of *Context-CDA*.

2. RELATED WORK

An LM may be deployed in a different setting than that for which it was intended, such as with or without a human intermediary for automated decision-making [38]. The primary causes of LM biases include inherent biases in the training data [1]. Bias mitigation techniques are categorized by the different stages of LM workflow. Pre-processing mitigation techniques focus on reducing bias and unfairness early in the dataset or model inputs. One of the early formalizations of this approach involves CDA, which has emerged as a prominent technique for mitigating biases in language models by introducing minimally perturbed examples that alter specific at-

tributes while preserving the overall semantics. The foundational work in CDA by [49] demonstrated its effectiveness in reducing gender bias in coreference resolution by swapping gendered terms in training data. [19] extended this to a more general framework by applying CDA for neural NLP tasks by generating matched sentence pairs through gendered word interventions. [52] proposed a CDA method tailored for morphologically rich languages to generate grammatically correct gender-swapped sentences.

As described by [44], the CDA procedure can be applied in a one-sided manner (using only the counterfactual sentence for further training) or a two-sided manner (incorporating both the original and counterfactual sentences in the training data). [20] introduce Counterfactual Data Substitution (CDS), a variant of CDA, where gendered terms in the training data are probabilistically replaced with their counterfactual counterparts, rather than duplicating each instance with a gender-swapped version. More recently, [45] proposed Polyjuice, a controllable counterfactual generation system based on GPT-2, enabling diverse perturbation types for training and evaluation. [25] introduced a neural demographic perturber trained on a large human-annotated dataset (PANDA), and demonstrate that augmenting data via demographic perturbations improves model fairness with minimal performance trade-offs. Both works highlight the utility of automated, fine-grained counterfactual generation for robust and fair NLP.

Recent advances in CDA span diverse domains, reflecting its growing relevance in addressing data imbalance and spurious correlations. In NLP, [41] introduce FairFlow, a model-based CDA method




Figure 2: StereoSet bias score for 5 LLM Models - Vanilla v/s CDA (Intrinsic bias). 50 indicates no bias.




Figure 3: StereoSet bias score for 5 LLM Models - Vanilla v/s Context-CDA.




Figure 4: CrowS-Pairs bias score for 5 LLM Models - Vanilla v/s CDA (Intrinsic bias). 50 indicates no bias.




Figure 5: CrowS-Pairs bias score for 5 LLM Models - Vanilla v/s Context-CDA.

that generates parallel counterfactuals without manual intervention, improving fairness while preserving fluency. Similarly, [35] employ self-imitation reinforcement learning for CDA, emphasizing both fairness and robustness under contextual shifts. In sentiment analysis, [46] propose polarity-reversing augmentations using pre-trained transformers to enhance aspect-based sentiment classification. Beyond text, CDA has also seen traction in graph neural networks. For instance, CAGAD [47] utilizes diffusion models to generate counterfactual anomalies in graphs, improving anomaly detection in node representations. In reinforcement learning, ACAMDA [37] recovers temporal causal structures to simulate realistic hypothetical scenarios for improved data efficiency. CAIAC [42] swaps causally irrelevant state components to generate robust offline learning transitions. A parallel effort to improve interpretability can be seen in FCE-UTD [17], which generates factor-level counterfactuals for causal explanations in Point-of-Interest (POI) recommendation. Together, these works underscore the versatility of CDA across modalities and its emerging role in mitigating bias.

Additionally, CDA has been integrated with other debiasing methods such as adversarial training [28] and calibration strategies [40] to improve robustness and generalization. However, concerns remain around the semantic validity and distributional shift introduced by counterfactuals, leading to research on controllable generation (e.g., using LMs or paraphrasing systems) and evaluation metrics to ensure linguistic and contextual coherence. Our work

builds upon this by extending the CDA framework to include context-rich, gender-flipped sentences generated by a large LM following the work of [10] for debiasing along with semantic entropy filtering to exclude ambiguous counterfactuals and improve corpus quality.

3. METHODS

The proposed method consists of three major steps: (1) Augmenting the context of CDA using large LMs, (2) Uncertainty-based filtering, and (3) Debiasing via fine-tuning on filtered counterfactuals. The overview of *Context-CDA* is illustrated in Figure 1.

CrowS-Pairs bias score for 5 LLM Models - Vanilla v/s CDA (Intrinsic bias). 50 indicates no bias.

3.1 Context-Aware CDA

For gender bias, vanilla CDA alters gender-related words, which can degrade language modeling ability while being effective at bias mitigation [8; 29]. One primary reason for this degradation is that the distribution of the corpus generated by vanilla CDA and then used for debiasing drifts from the distribution of the pre-trained corpus. When debiasing via fine-tuning the counterfactual generation, the n-gram distribution of the target LM is altered, resulting in the degradation of language modeling ability [7]. To address the limitation, we introduce *Context-CDA*, a prompting-based method that advances traditional CDA by prompting a larger LM to gener-




Figure 6: LMS score (↑) for 5 LLM Models - Vanilla v/s CDA.




Figure 8: ICAT score (↑) for 5 LLM Models - Vanilla v/s CDA.

ate contextually richer counterfactuals. Rather than merely flipping gender-related words, *Context-CDA* refines the sentences to preserve their original meaning while providing a more natural and diverse context aligned with the LM’s pre-training distribution. This approach to context-aware data augmentation mitigates the severe distortions often caused by traditional CDA methods, where oversimplified examples fail to account for the social context of altered sensitive attributes. As a result, our method enables a more thorough debiasing process, promoting fairness while preserving the LMs’ modeling capabilities. Specifically, after generating counterfactual examples using CDA $\tilde{x}_i = \text{CDA}(x_i)$, we design a system prompt and an instruction prompt (illustrated in Figure 1) to ask a Llama-3-8b-Instruct [9] model to rephrase \tilde{x}_i such that it contains more context. The resulting generation is denoted as \tilde{x}_i^c .

3.2 Uncertainty-Based Filtering

With the generated counterfactuals, another challenge is that text generated by the larger LMs can be overly complex or noisy for the target smaller LMs to be debiased, hindering learning effectiveness. To further enhance the quality (e.g., mitigating hallucination) of the generated counterfactuals for the target LMs, we propose filtering out the generations that exhibit the greatest uncertainty by the target LMs. Uncertainty in the data has been shown to be an informative signal for algorithmic bias [39; 33] and hallucination in large LMs [6]. Specifically, we use a filtering process based on semantic entropy following the methodology in [15], which advances

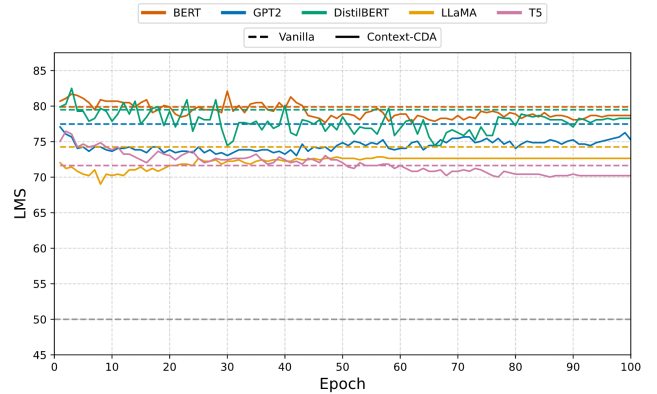


Figure 7: LMS score (↑) for 5 LLM Models - Vanilla v/s *Context-CDA*.




Figure 9: ICAT score (↑) for 5 LLM Models - Vanilla v/s *Context-CDA*.

other uncertainty estimation approaches like likelihoods as it incorporates linguistic invariances created by shared meanings. For this process, we first sample multiple rephrased output sequences from a language model for each of our generated counterfactuals. These sequences are then clustered into semantic equivalence classes using a bidirectional entailment test following the methodology in [15], where two sequences are considered equivalent if they logically imply each other within context. Finally, the probabilities of sequences in each cluster are summed, and semantic entropy is calculated over these meaning-level probabilities to measure uncertainty over meanings rather than just surface forms. The semantic entropy (SE) for each counterfactual can be calculated as follows:

$$SE(\tilde{x}_i^c) \approx -|C|^{-1} \sum_{j=1}^{|C|} \log p(C_j | \tilde{x}_i^c), \quad (1)$$

where C denotes the number of samples generated by the larger LM. We calculate the semantic entropy for each generated counterfactual \tilde{x}_i^c using the target LMs and filter out the top k -percent (e.g., 30%) of sentences with the highest entropy. This ensures that overly complex or noisy sentences, which the target LMs may struggle to learn, are removed from the corpus, leaving only the most useful examples for debiasing during the fine-tuning.

3.3 Debiasing via Fine-tuning on Filtered Context-CDA

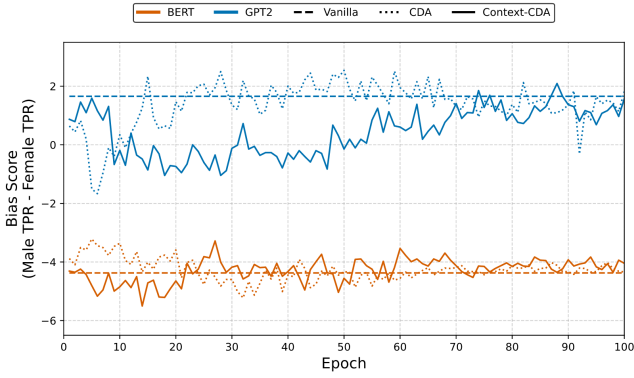


Figure 10: BiasBios score for BERT and GPT-2 (Extrinsic bias). 0 indicates no bias.

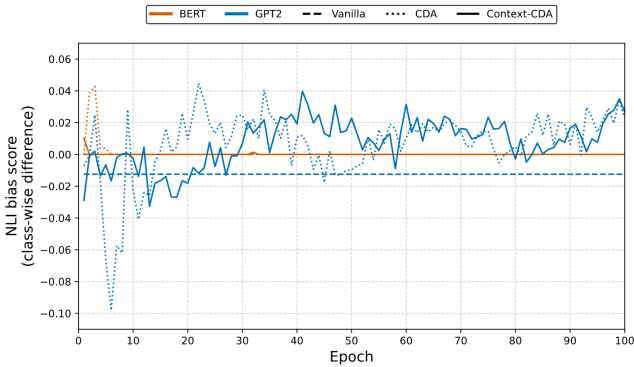


Figure 12: NLI-Bias score for BERT and GPT-2 (Extrinsic bias). 0 indicates no bias.

We use the contextually rich counterfactual data samples obtained after augmentation and filtering to debias the target small LMs [10] like BERT and GPT-2. The key idea is to introduce alternative versions of the input data where specific attributes (e.g., gender) are modified without changing the underlying meaning via augmented context. During fine-tuning, this augmentation forces the model to learn representations that are invariant to these modifications, which helps mitigate biased correlations in the data.

While our evaluation focuses on encoder models like BERT, this kind of representation debiasing is relevant to generative systems too. This is largely because BERT-like transformer encoders serve as foundational components in state-of-the-art encoder-decoder models (e.g., T5, BART, mBART) which are widely used for generation tasks such as machine translation and summarization. Moreover, contextualized encoder representations are commonly employed as frozen or fine-tuned backbones in modular and retrieval-augmented generation pipelines [16]. As a result, debiasing these shared representations can directly benefit a broad range of downstream generative applications.

Moreover, to demonstrate generalizability of *Context-CDA*, we evaluate across diverse architectures including encoder-only (BERT, DistilBERT), encoder-decoder (T5), and decoder-only (GPT-2, Llama-3.2-1B) models, as detailed in Section 4. In summary, our *Context-CDA* method aims to balance debiasing and language modeling performance by generating augmented sentences that maintain linguistic coherence while reducing bias via fine-tuning. This ensures




Figure 11: STS-B bias score for BERT and GPT-2 (Extrinsic bias). 0 indicates no bias.




Figure 13: QNLI score (\uparrow) for BERT and GPT-2 (Downstream task).

that the target model learns from a corpus that is both diverse and aligned with natural language patterns.

4. EXPERIMENTS

We conduct the following experiments to validate the effectiveness of our proposed approach: (1) Evaluating the debiasing performance on intrinsic bias and language modeling capabilities, (2) Evaluating the debiasing performance on extrinsic bias, (3) Evaluating the debiasing performance on various downstream tasks, and (4) Analyzing next-token distribution to study predicted output token shifts after debiasing. All experiments are conducted across five diverse model architectures spanning encoder-only, encoder-decoder, and decoder-only designs to comprehensively validate *Context-CDA*'s robustness and generalizability.

4.1 Experimental Setup

Datasets: To implement gender bias mitigation, we use the news-commentary dataset [36] as our debiasing corpus following [14] with the gender words list from [50]. For **intrinsic bias**, we use gender samples from two benchmark datasets: (1) StereoSet [23] tests whether models can complete sentences without reinforcing harmful stereotypes, while still maintaining language fluency. (2) CrowS-Pairs [24] tests social biases in LMs by presenting sentence pairs, where one reflects a stereotype and the other does not. For **extrinsic bias**, we use datasets (1) BiasBios (Bias in Bios)




Figure 14: RTE score (↑) for BERT and GPT-2 (Downstream task).

[4], (2) STS-B (Semantic Textual Similarity Benchmark) [2] and (3) NLI-Bias (Natural Language Inference-Bias) [4]. For evaluating model performances on downstream tasks like sentiment analysis, entailment, and question-answering, we use the following datasets from the GLUE (General Language Understanding Evaluation) benchmark [43], (1) SST-2 (Stanford Sentiment Treebank-2) [34], (2) RTE (Recognizing Textual Entailment) [3], and (3) QNLI (Question-answering Natural Language Inference) [43] derived from the Stanford Question Answering Dataset (SQuAD).

Base LMs: Our evaluation covers five diverse LMs spanning different architectures. This includes two encoder-only models (BERT [5] and DistilBERT [30]), one encoder-decoder model (T5 [27]), and two causal decoder-only models (GPT-2 [26] and Llama-3.2-1B [13]) for a comprehensive evaluation. We evaluate BERT and GPT-2 on all intrinsic bias, extrinsic bias, and downstream task performance metrics. DistilBERT, T5, and Llama-3.2-1B are evaluated solely on intrinsic bias metrics. This selection enables us to validate that *Context-CDA* is model-agnostic and effective across encoder-only, encoder-decoder, and decoder-only architectures.

Baselines: We compare the Vanilla model, which is the pre-trained LM without debiasing, and LMs debiased using traditional CDA as the baseline for all models including BERT, GPT-2, DistilBERT, T5 and Llama-3.2-1B. Additionally, we consider several baselines for BERT and GPT-2, respectively. For BERT, we use the following methods as the baseline: (1) MABEL [11], an intermediate pre-training approach for mitigating gender bias in contextualized representation; (2) INLP [28], mitigating gender bias in word embeddings by removing information from neural representations; (3) SelfDebias [31], a decoding algorithm that, given only a textual description of the undesired behavior, reduces the probability of an LM producing problematic text; and (4) SENT-DEBIAS [18], which reduces gender bias in sentence-level representations. For GPT-2, we use the following methods as baselines: (1) SelfDebias [31], a decoding algorithm that, given only a textual description of the undesired behavior, reduces the probability of an LM producing problematic text; (2) SENT-DEBIAS [18], which reduces gender bias in sentence-level representations, and (3) wiki-debiased [48], a baseline with GPT-2 as the base LM that uses parameter-efficient methods to fine-tune GPT-2 using WikiText2 [21].

Metrics: To evaluate debiasing performance, we distinguish between intrinsic bias and extrinsic bias. Intrinsic bias refers to the stereotypical associations encoded directly within a language model’s representations or probabilities, independent of downstream tasks.




Figure 15: SST-2 score (↑) for BERT and GPT-2 (Downstream task).

It is commonly measured through benchmarks such as StereoSet and CrowS-Pairs, which assess whether models favor stereotypical over anti-stereotypical continuations or sentence pairs. Extrinsic bias captures disparities that arise when models are applied to downstream tasks. It reflects whether representational biases affect task performance, for instance, differences in prediction accuracy or true positive rates across demographic groups. We measure extrinsic bias using datasets such as BiasBios, STS-B, and NLI-Bias.

For intrinsic bias, we report four standard metrics from the StereoSet [23] and CrowS-Pairs [24] benchmarks: (1) Stereotype Score (SS) calculates the percentage of examples in which a model prefers a stereotypical association over an anti-stereotypical association (a score of 50 indicates no bias and a score above and below 50 indicates preference for stereotypical or anti-stereotypical associations); (2) Language Modeling Score (LMS) calculates the percentage of instances in which an LM prefers meaningful over meaningless association (a higher score indicates better language performance); (3) Idealized CAT Score (ICAT) describes the comprehensive performance of the model by combining SS and LMS into one score (a higher score indicates better performance on debiasing while maintaining language fluency); and (4) CrowS-Pairs Score (CS) calculates the percentage of instances where the model assigns a higher probability to the stereotypical sentence over the anti-stereotypical one (a score closer to 50 indicates less bias, while scores above or below 50 suggest a preference toward biased associations).

For extrinsic bias, we report the following three metrics: (1) NLI-Bias measures the class-wise difference between male and female samples in the NLI-Bias dataset, with scores closer to 0 indicating less bias; (2) BiasBios measures the difference in true positive rate between male and female samples in BiasBios dataset; and (3) STS-B flips the gendered words in the sentence pairs and then calculates semantic similarity in original and gender-flipped sentence pairs to measure gender bias. The bias score is calculated as the difference in prediction accuracy between male and female groups, with scores closer to 0 indicating less bias. For performance on downstream tasks, we use accuracy as a metric for datasets (1) QNLI, (2) RTE, and (3) SST-2.

We set the uncertainty filtering threshold k to 30% as it gives the best performance based on our debiasing corpus as shown in the ablation study in Appendix A. For each language model, including BERT, GPT-2, DistilBERT, T5, and Llama-3.2-1B, we perform debiasing at each epoch to evaluate the intrinsic bias and lan-

Debiasing Technique	BERT				GPT-2			
	SS	LMS (\uparrow)	ICAT (\uparrow)	CS	SS	LMS (\uparrow)	ICAT (\uparrow)	CS
MABEL	47.28	51.65	48.84	<u>52.29</u>	-	-	-	-
INLP	<u>49.16</u>	50.25	49.41	<u>55.73</u>	-	-	-	-
wiki-debiased	-	-	-	-	60.40	<u>91.01</u>	72.08	56.49
SelfDebias	59.34	<u>84.20</u>	<u>68.47</u>	<u>52.29</u>	<u>56.05</u>	87.43	<u>73.18</u>	56.11
SENT-DEBIAS	59.37	84.09	68.33	<u>52.29</u>	60.84	89.07	69.76	56.11
Vanilla	59.95	79.87	63.96	58.01	63.17	77.46	57.04	51.52
CDA	58.55	68.41	56.71	54.96	57.34	69.61	59.39	<u>50.01</u>
<i>Context-CDA</i>	<u>57.75</u>	78.67	66.48	53.43	56.13	75.25	66.01	50.76

Table 1: Intrinsic bias evaluation for BERT and GPT-2 comparing baselines. Underlined values indicate the best performance.

guage modeling performance using the StereoSet and CrowS-Pairs benchmarks. Since our method fine-tunes models, we compare the vanilla, CDA, and *Context-CDA* models at each epoch of debiasing. For BERT and GPT-2, we additionally fine-tune models at each epoch to measure extrinsic bias and downstream task performance across datasets. This ensures a comprehensive assessment of debiasing effectiveness across both representation-level and task-level biases. This multi-faceted evaluation enables us to assess not only whether bias is reduced but also whether language modeling capability and downstream task performance are preserved.

4.2 Post Debiasing Performance on Intrinsic Bias Scores and Language Modeling

4.2.1 Multi-Model Evaluation: Demonstrating Robustness and Generalizability

To validate the robustness and model-agnostic nature of *Context-CDA*, we evaluate intrinsic bias on all five LMs.

Intrinsic Bias. To measure intrinsic bias, we conduct evaluations using the StereoSet (SS) and CrowS-Pairs (CS) benchmark at every debiasing epoch across all model architectures - BERT, DistilBERT, GPT-2, Llama-3.2-1B, and T5. We present detailed results comparing the performance of all models based on the SS score for CDA in Fig. 2 and *Context-CDA* in Fig. 3. We also present similar detailed results based on the CS score for CDA in Fig. 4 and *Context-CDA* in Fig. 5. Here, scores closer to 50 indicate less bias with a score of 50 indicating no bias. We observe that the vanilla BERT and DistilBERT generally start with the highest bias scores before debiasing begins. As debiasing progresses across epochs, both encoder models show consistent debiasing patterns and the SS and CS scores gradually approaches 50, indicating a reduction in bias in the models. We also observe that towards the end of training epochs, the bias scores stabilize and we conclude that further fine-tuning is not required. Towards epochs 75-85, *Context-CDA* starts outperforming CDA and achieves a stable bias score better than CDA. This consistency validates *Context-CDA* for encoder-only architectures. In GPT-2, CS bias score for *Context-CDA* and traditional CDA reaches around 50 (no bias) early on and oscillates about the same mean bias score as fine-tuning progresses. In Llama-3.2-1B, epochs 0 to 25 show a sustained decline in SS score reaching a plateau that indicates convergence without overfitting. The stable plateau throughout training demonstrates that further fine-tuning does not degrade or improve the debiasing performance. Both generative language models achieve comparable or superior results to their non-augmented baselines, supporting our claim that *Context-CDA* is effective for generative systems as well. T5, as an encoder-decoder model, shows particularly strong performance, achieving a balanced SS score and a significant improvement in CS score. Thus, compared to traditional CDA, *Context-CDA* shows

consistent improvement across all intrinsic metrics.

Language Modeling Ability. To measure language modeling performance, we evaluate the Language Modeling Score (LMS) and Idealized CAT Score (ICAT) scores at every debiasing epoch across all model architectures - BERT, DistilBERT, GPT-2, Llama-3.2-1B, and T5. We present detailed results comparing the performance of all models based on the LMS score for CDA in Fig. 6 and *Context-CDA* in Fig. 7. We also present similar detailed results based on the ICAT score for CDA in Fig. 8 and *Context-CDA* in Fig. 9. Notably, across all models, *Context-CDA* consistently improves both LMS and ICAT scores compared to CDA. *Context-CDA* consistently outperforms CDA and achieves LMS scores as well as Vanilla models, while outperforming both Vanilla and CDA baselines in ICAT scores. This is a strong indicator of *Context-CDA*'s superior language modeling performance compared to CDA, reinforcing that *Context-CDA* achieves better linguistic understanding and better preserves language modeling capability while reducing bias. We attribute this improvement to the greater diversity and naturalness of the *Context-CDA* corpus compared to traditional CDA, which enables models to retain stronger linguistic representations after debiasing. Overall, *Context-CDA* achieves effective bias mitigation while simultaneously improving language modeling performance.

4.2.2 Convergence Patterns

The per-epoch analysis reveals important training dynamics: (1) **Convergence Timing:** Across all five models, debiasing performance stabilizes at epochs 75-85, indicating the method reaches optimal performance earlier in training. (2) **Stability Without Overfitting:** The plateau in bias scores and sustained or improved scores demonstrate that models do not overfit to the debiasing corpus; instead, they learn stable, debiased representations. (3) **Model-Agnostic Pattern:** The identical convergence behavior across BERT, DistilBERT, GPT-2, Llama-3.2-1B, and T5 validates that this stability is not an artifact of a single architecture but a general property of *Context-CDA*. (4) **Training Efficiency:** The early stabilization suggests that practitioners can reduce training epochs, making the method computationally efficient. These insights strengthen confidence in the method's robustness and generalizability.

4.2.3 Comparison with Prior Debiasing Methods

We also compare the bias scores of fine-tuned BERT and GPT-2 against other debiasing baselines in Table 1 for an overall comparison. While methods such as SelfDebias and SENT-DEBIAS achieve higher scores on isolated metrics (e.g., ICAT), they often rely on inference-time interventions or post-hoc representation manipulation, which: (1) are model-specific or task-specific, limiting generalizability; (2) do not demonstrate consistent performance




Figure 16: Top-20 male token distributions for GPT2.




Figure 17: Top-20 female token distributions for GPT2.

across both intrinsic and extrinsic metrics; and (3) lack validation on diverse model architectures. In contrast, our extended evaluation demonstrates that *Context-CDA* achieves balanced, consistent performance across five distinct architectures, maintaining stable debiasing while preserving language modeling capability across both encoder-only and generative models. Wiki-debiased, in contrast, leverages curated Wikipedia data, making it domain-limited, whereas *Context-CDA* can be applied to arbitrary corpora. Consequently, *Context-CDA* is model-agnostic, requires no architecture modifications, and integrates seamlessly into the fine-tuning pipeline, making it a practical and effective debiasing solution across model architectures. The consistency of results across five model types provides strong empirical evidence of robustness, addressing concerns about limited generalization from single-model evaluation. It strikes a strong balance by offering the best trade-off between bias mitigation and language modeling performance, outperforming standard CDA across all intrinsic bias metrics and achieving comparable or superior results to other baselines, all without sacrificing fluency.

4.3 Post Debiasing Performance on Extrinsic Bias and Downstream Tasks

For evaluating performance on extrinsic bias metrics, we further fine-tune our models - BERT and GPT-2 on STS-B, NLI-Bias and BiasBios tasks and then evaluate the extrinsic bias scores on the fine-tuned models. Here, scores closer to 0 indicate low bias. For BERT, we find that for STS-B (Fig. 11), *Context-CDA* performs better than CDA, whereas for NLI-Bias and BiasBios (Fig. 12 and Fig. 10 respectively), *Context-CDA* performs as good as CDA indicating that *Context-CDA* is robust for even extrinsic bias evaluation metrics. For GPT-2, *Context-CDA* performs as well as CDA for BiasBios, STS-B, and NLI-Bias (Figs. 10, 11, and 12 respectively), further validating its effectiveness in reducing extrinsic bias across model architectures. For evaluating performance on various

downstream language understanding tasks, we further fine-tune our debiased models on QNLI, RTE, and SST-2 and then evaluate the fine-tuned models for accuracy. Higher score indicates higher accuracy. For QNLI (Fig. 13), *Context-CDA* performs better than CDA whereas for RTE and SST-2 (Fig. 14 and Fig. 15), *Context-CDA* performs slightly better or as good as CDA, indicating that models fine-tuned with *Context-CDA* can perform better than CDA even in various downstream tasks.

4.4 Next-Token Distribution

To gain deeper insights into the debiasing performance, we compare the vanilla LM, CDA, and *Context-CDA* by examining next-token distributions in gender-related contexts. We focus on the GPT-2 model with the Stereotype Score (SS) closest to 50 after CDA debiasing, allowing us to investigate how the debiasing process affects token distribution and explore its impact on model predictions. We use sentences from StereoSet that specifically evaluate gender bias. Each sentence is split into two parts: the portion before the BLANK, referred to as the context, and the portion after it. The context is fed into GPT-2, and we compute the logits for the next token, extracting scores for tokens in a predefined male-female mapping set (approximately 200 words) [50]. After applying softmax, we obtain probabilities for each word and identify the top-20 tokens. We then count how many tokens are male-related and how many are female-related.

Figures 16 and 17 illustrate the frequency of male and female tokens, respectively, among the top-20 predictions. For male tokens, the vanilla LM peaks at 7 male-related tokens, while CDA and *Context-CDA* shift the peak to 8. Importantly, *Context-CDA* yields a more balanced distribution across contexts, reducing skew toward male-associated tokens. Similarly, for female tokens, the vanilla LM peaks at 13 female-related tokens, while CDA and *Context-CDA* shift the peak toward 12. Again, *Context-CDA* produces a more even distribution, avoiding over-concentration on specific female-

related tokens. Together, these results demonstrate that both CDA and *Context-CDA* adjust the token distribution away from the stronger bias seen in the Vanilla model. However, *Context-CDA* is more effective at spreading token probabilities evenly across male and female categories. This not only mitigates over-reliance on gender-specific terms but also promotes linguistic diversity, thereby improving robustness in predictions and reflecting a more balanced language modeling ability.

5. LIMITATIONS

While our method *Context-CDA* demonstrates improvements over traditional CDA in mitigating gender bias without compromising language modeling performance, several limitations remain. First, using large LMs for generating context-rich augmentations may introduce computational and environmental costs, which may hinder scalability and accessibility in resource-constrained settings. Second, although semantic entropy filtering improves corpus quality by excluding uncertain examples, it may also remove valuable complex counterfactuals; future work could explore adaptive or multi-criteria filtering strategies that balance quality, diversity, and training stability. Third, our study focuses primarily on binary gender counterfactuals. Extending this framework to non-binary and intersectional identities, as well as other sensitive attributes (e.g., race, religion, or profession), is an important next step. Fourth, because large LMs and target smaller models may differ in distributional characteristics, alignment mechanisms such as domain-adaptive filtering, grounding mechanisms such as fact-checking modules, and human-in-the-loop validation could further enhance contextual reliability. Finally, expanding this framework to include domain-specific bias detection or multilingual or multimodal extensions would enhance robustness and fairness across broader applications.

6. CONCLUSION

This work presents an effective gender debiasing method that maintains competitive performance in downstream tasks while reducing bias in LMs. Building on classic CDA, which effectively mitigates bias but often weakens language modeling capabilities, our proposed method, *Context-CDA*, enhances the debiasing corpus by leveraging large LMs to generate enriched context. This augmentation minimizes the discrepancy between the debiasing corpus and the original pre-training data, ensuring better alignment and also preserving linguistic fluency. Furthermore, we incorporate semantic entropy filtering to remove uncertain content, improving the overall quality of the generated corpus. Comprehensive evaluation across five diverse model architectures demonstrates that *Context-CDA* is truly model-agnostic, achieving robust and consistent debiasing performance across both discriminative and generative systems. Our method not only mitigates bias effectively but also enhances language modeling performance. As LMs continue to evolve, integrating more sophisticated debiasing techniques will be crucial for building more equitable and more robust AI systems.

7. REFERENCES

- [1] Baumann et al. Bias on demand: a modelling framework that generates synthetic data with bias. In *Proceedings of the 2023 ACM Conference on FAcCT*, pages 1002–1013, 2023.
- [2] Cer et al. Semeval-2017 task 1: Semantic textual similarity-multilingual and cross-lingual focused evaluation. *arXiv preprint arXiv:1708.00055*, 2017.
- [3] Dagan et al. The pascal recognising textual entailment challenge. In *ML challenges workshop*. Springer, 2005.
- [4] De-Arteaga et al. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on FAcCT*, pages 120–128, 2019.
- [5] Devlin et al. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of ACL*, 2019.
- [6] Farquhar et al. Detecting hallucinations in large language models using semantic entropy. *Nature*, 630(8017), 2024.
- [7] Fatemi et al. Improving gender fairness of pre-trained language models without catastrophic forgetting. 2021.
- [8] Gallegos et al. Bias and fairness in large language models: A survey. *Computational Linguistics*, 50(3):1097–1179, 2024.
- [9] Grattafiori et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [10] Han et al. Chatgpt based data augmentation for improved parameter-efficient debiasing of llms. 2024.
- [11] He et al. Mabel: Attenuating gender bias using textual entailment data. *arXiv preprint arXiv:2210.14975*, 2022.
- [12] Huang et al. Counterfactually-augmented snli training data does not yield better generalization than unaugmented data. *arXiv preprint arXiv:2010.04762*, 2020.
- [13] Hugging Face. meta-llama/llama-3.2-1b. <https://huggingface.co/meta-llama/Llama-3.2-1B>, 2024.
- [14] Kaneko et al. Debiasing isn’t enough!—on the effectiveness of debiasing mlms and their social biases in downstream tasks. *arXiv preprint arXiv:2210.02938*, 2022.
- [15] Kuhnet et al. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. *arXiv preprint arXiv:2302.09664*, 2023.
- [16] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W. tau Yih, T. Rocktäschel, S. Riedel, and D. Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks, 2021.
- [17] Li et al. Beyond relevance: Factor-level causal explanation for user travel decisions with counterfactual data augmentation. *ACM Transactions on Information Systems*, 2024.
- [18] Liang et al. Towards debiasing sentence representations. In *Proceedings of 58th Annual Meeting of the ACL*, July 2020.
- [19] Lu et al. Gender bias in neural natural language processing. *Logic, language, and security: essays dedicated to Andre Scedrov on the occasion of his 65th birthday*, 2020.
- [20] Maudslay et al. It’s all in the name: mitigating gender bias with name-based counterfactual data substitution. *arXiv preprint arXiv:1909.00871*, 2019.
- [21] Merity et al. Pointer sentinel mixture models, 2016.
- [22] Miller et al. Parlai: A dialog research software platform. *arXiv preprint arXiv:1705.06476*, 2017.
- [23] Nadeem et al. Stereoset: Measuring stereotypical bias in pre-trained language models. *Preprint arXiv:2004.09456*, 2020.

- [24] Nangia et al. Crows-pairs: A challenge dataset for measuring social biases in masked language models. *arXiv preprint arXiv:2010.00133*, 2020.
- [25] Qian et al. Perturbation augmentation for fairer nlp. *arXiv preprint arXiv:2205.12586*, 2022.
- [26] Radford et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [27] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer, 2023.
- [28] Ravfogel et al. Null it out: Guarding protected attributes by iterative nullspace projection. *arXiv preprint arXiv:2004.07667*, 2020.
- [29] Raza et al. Mbias: Mitigating bias in large language models while retaining context. *Preprint arXiv:2405.11290*, 2024.
- [30] V. Sanh, L. Debut, J. Chaumond, and T. Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, 2020.
- [31] Schick et al. Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp, 2021.
- [32] Sheng et al. The woman worked as a babysitter: On biases in language generation. *arXiv preprint arXiv:1909.01326*, 2019.
- [33] Singh et al. Fairness in ranking under uncertainty. *Advances in Neural Information Processing Systems*, 2021.
- [34] Socher et al. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- [35] K. Sreedhar, T. Kavva, J. Prasad, and V. Varshini. A novel metric-based counterfactual data augmentation with self-imitation reinforcement learning (sil). *International Journal of Advanced Computer Science & Applications*, 16(1), 2025.
- [36] Statmt. Statistical and neural machine translation, 2023.
- [37] Sun et al. Acamda: improving data efficiency in reinforcement learning through guided counterfactual data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 15193–15201, 2024.
- [38] Suresh et al. A framework for understanding sources of harm throughout the machine learning life cycle. In *Proceedings of the 1st ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, pages 1–9, 2021.
- [39] Tahir et al. Fairness through aleatoric uncertainty. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 2372–2381, 2023.
- [40] Tan et al. Assessing social and intersectional biases in contextualized word representations. *Advances in neural information processing systems*, 32, 2019.
- [41] Tokpo et al. Fairflow: An automated approach to model-based counterfactual data augmentation for nlp. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 160–176. Springer, 2024.
- [42] Urpí et al. Causal action influence aware counterfactual data augmentation. *arXiv preprint arXiv:2405.18917*, 2024.
- [43] Wang et al. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- [44] Webster et al. Measuring and reducing gendered correlations in pre-trained models. *Preprint arXiv:2010.06032*, 2020.
- [45] Wu et al. Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models. *arXiv preprint arXiv:2101.00288*, 2021.
- [46] Wu et al. A novel counterfactual data augmentation method for aspect-based sentiment analysis. In *Asian Conference on Machine Learning*, pages 1479–1493. PMLR, 2024.
- [47] Xiao et al. Counterfactual data augmentation with denoising diffusion for graph anomaly detection. *IEEE Transactions on Computational Social Systems*, 2024.
- [48] Xie et al. An empirical analysis of parameter-efficient methods for debiasing pre-trained language models, 2023.
- [49] Zhao et al. Gender bias in coreference resolution: Evaluation and debiasing methods. *Preprint arXiv:1804.06876*, 2018.
- [50] Zhao et al. Learning gender-neutral word embeddings. *arXiv preprint arXiv:1809.01496*, 2018.
- [51] Zhu et al. Multilingual machine translation with llms: Empirical results and analysis. *arXiv preprint arXiv:2304.04675*, 2023.
- [52] Zmigrod et al. Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology. *arXiv preprint arXiv:1906.04571*, 2019.

APPENDIX

A. ABLATION STUDY ON UNCERTAINTY THRESHOLD

To validate the robustness of our proposed *Context-CDA* approach, we conduct an ablation study analyzing the effect of varying the semantic entropy filtering threshold on debiasing and language modeling performance. In our main experiments, we used a default threshold of 30%, removing the top 30% of counterfactuals with the highest semantic entropy. This study evaluates how different thresholds (20% and 40%) influence the effectiveness of the model.

Impact on debiasing performance: Figures 18–21 illustrate the effect of semantic entropy thresholds on intrinsic bias metrics such as Stereotype Score (SS) and CrowS-Pairs Score (CS) for both BERT and GPT-2. At a 20% threshold (i.e., more lenient filtering), the model retains more counterfactuals, including those with moderate uncertainty. This sometimes leads to insufficient bias removal, particularly visible in the slightly higher CS values in Fig. 19. Conversely, at a 40% threshold (i.e., more aggressive filtering), while bias mitigation improves initially due to the exclusion of noisier samples, over-filtering may reduce the diversity of the counterfactual corpus, potentially limiting the coverage of gender-related variations and reducing generalization. Overall, the 30% threshold achieves the best balance: it significantly reduces gender bias and outperforms CDA while avoiding the potential drawbacks of both under- and over-filtering as observed in Fig. 19, 20 and 21.




Figure 18: BERT SS Score for SE thresholds 20, 30 and 40.




Figure 19: BERT CS Score for SE thresholds 20, 30 and 40.




Figure 20: GPT-2 SS Score SE thresholds 20, 30 and 40.




Figure 21: GPT-2 CS Score for SE thresholds 20, 30 and 40.




Figure 22: BERT LMS Score for SE thresholds 20, 30 and 40.




Figure 23: BERT ICAT Score for SE thresholds 20, 30 and 40.




Figure 24: GPT-2 LMS Score for SE thresholds 20, 30 and 40.




Figure 25: GPT-2 ICAT Score for SE thresholds 20, 30 and 40.

Hence, we choose 30% as the optimal setting for entropy-based filtering in the main experiments. However, the higher SS scores in Fig. 18 indicate the need for a more fine-tuned evaluation of this threshold to arrive at a balanced and optimal value.

Impact on Language Modeling Performance: Figures 22–25 report the Language Modeling Score (LMS) and ICAT score under different entropy thresholds-20, 30 and 40. For both BERT and GPT-2, a 20% threshold tends to retain too many noisy samples, leading to slightly reduced fluency and coherence, as seen in lower LMS scores in Fig. 22. On the other hand, a 40% threshold, while helping eliminate uncertain samples, may discard too many useful and contextually rich augmentations. In contrast, the 30% threshold yields higher LMS and ICAT scores, indicating better preservation of language fluency and semantic understanding. These results reinforce that filtering at this level helps strike an optimal trade-off between reducing uncertainty and retaining the linguistic richness

of the training corpus. However, these thresholds can be further refined to reach a more balanced threshold.