E-ISSN: 1115 – 8492 P-ISSN: 3027 – 0693
MPRI-JNSAR | INTERNATIONAL JOURNAL OF NATURE AND SCIENCE ADVANCE RES.
VOL. 10 NO. 9 OCTOBER-2025

MITIGATING SOCIAL AND CULTURAL BIASES IN PRE-TRAINED LANGUAGE MODELS: A LINGUISTICALLY INFORMED DEBIASING FRAMEWORK

ISHOLA, P.E.; OLOFIN, B.B.; ISHOLA, A.K.; ADEOLA, P.; & UMOREN, U.M.
Department of Computing, Anchor University Lagos, Nigeria
Corresponding Author:
patience.ishola@yabatech.edu.ng
DOI: https://doi.org/10.70382/mejnsar.v10i9.066

Abstract

Pre-trained language models have demonstrated exceptional performance across diverse natural language processing tasks; however, they often inherit and amplify the social and cultural biases present in their training data. These biases can lead to unfair and discriminatory outcomes, limiting the ethical deployment of AI systems. This study introduces a linguistically informed debiasing framework designed to mitigate such biases in publicly available pre-trained models. By leveraging deep linguistic insights—including syntactic, semantic, and pragmatic structures—our approach systematically identifies and reduces biased representations without compromising model efficacy. The proposed methodology integrates counterfactual data augmentation, culturally balanced corpora, and targeted fine-tuning strategies rooted in linguistic theory. Experimental results demonstrate that our framework effectively diminishes bias while preserving model performance, promoting fairness and inclusivity in real-world NLP applications.

Keywords: Social Biases, Cultural Biases, Pre-trained Language Models, Linguistically Informed Debiasing, Counterfactual Data Augmentation, Fairness

Introduction

The existing Pre-trained language models (PLMs) have become fundamentals in natural language processing (NLP), showing remarkable performance across various mission critical implementation. However, these models often imbibe and resonates social and cultural biases apparent in their training dataset, which resulted to

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

concerns about fairness and ethical implementation. Related studies have defined that PLMs can perpetuate stereotypes and discriminatory language, affecting marginalised groups or low level langugage conformity and undermining the reliability and versatility of Artificial Intelligent systems in real-world platforms.

To solve the prevailing concerns, related works have explored various debiasing benchmark and techniques. For instance, (Anantaprayoon et al., 2025) proposed an intent-aware self-correction and content basis mechanism that ingrains explicit debiasing prompts and chain-of-thought reasoning to mitigate biases in large language models. Again, (Chen et al., 2025) elucidates the Fairness Stamp (FAST) method, which bounded specific layers within PLMs to alleviate social biases while preserving overall model functionality, representation and performance. These methodologies demonstrate the potential of incorporating linguistic commonalities and insights to the targeted interventions with the purpose of enhanced model fairness.

Checking the richness of hindsight to the prevailing advancements, challenges surfaced in effectively prototyping, identifying and mitigating biases across multiple languages and social-cultural contexts. (Mitchell, 2025) exemplified the essence of culturally informed datasets, such as SHADES, to evaluate and address AI biases beyond English, focusing on the tagged low level language. This pins the need for linguistically informed frameworks that consider syntactic and semantic variations across multiple languages to ensure ethical and nonnarrowed systems with broader scalability and extensibility.

In this research work, we explored a linguistically informed debiasing framework that centred on mitigating social and cultural biases in PLMs. The cross-linguistic insights and incorporating culturally sensitive evaluation metrics are employed with resilient approach that seek to enhance the fairness and inclusivity of language models deployed in diverse realworld conundrum.

RELATED WORKS

Pre-trained language models (PLMs) have changed the landscape of natural language processing (NLP) by achieving remarkable improvement and surpassing capability across various problem-solving processes. However, a noticeable social and cultural biases raise discussion about decision taking on mission critical systems. Recent research works have explored several methodologies to identify and mitigate these biases, and pointed the need for linguistically informed frameworks. provides an intent-aware self-correction mechanism that makes good use of explicit debiasing prompts and chain-of-thought reasoning techniques to reduce biases in an elucidated large language models.

(Chen et al, 2025) postulates the Fairness Stamp (FAST) method, which makes specific layers within PLMs to mitigate social biases while preserving overall model stability and

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

performance. (Xu et al, 2025) implemented 'FaIRMaker', an fully automated and model-independent framework that uses an auto-search and refinement paradigm to adaptively generate considerable Fairwords, which efficiently and effectively mitigating social prevalent issues such as gender bias while preserving the overall model core integrity.
(Liu et al, 2025) makes of Social Bias Neurons and promoted Integrated Gap Gradients (IG^2) to accurately pinpoint various units in a language model which often is pinpointed and attributed to undesirable behaviours. Their Bias Neuron Suppression (BNS) technique effectively mitigates social biases by targeting a measured specific neurons. make sure that you uses machine unlearning techniques, such as Partitioned Contrastive Gradient Unlearning (PCGU) and Negation via Task Vector, to mitigate social biases in a specific language models.
(Dai et al, 2024) released Continuous Prompts Adjustment Debiasing (CPAD) model, which generates continuous token lists from the entire vocabulary space to assuage the gap between outputs and targets in a community fairness dominated learning processes.
(Yu et al, 2025) addressed gender biases in PLMs by enhancing fairness through the integration of LLM-generated sentences. (Dawkins et al., 2024) promulgates projective methods for mitigating gender bias in a definite PLMs. (Mitchell 2025) showed the importance of culturally informed datasets, such as SHADES, to evaluate and address social and cultural biases beyond English with special consderation to low-level or low-valued languages. (Dong et al, 2023) introduced a tool Co^2PT, which is a debias-while-prompt tuning method that mitigates biases via counterfactual contrastive prompt tuning on downstream activities.
(Li et al., 2024) shows fairBERTs, a general framework for learning fair fine-tuned BERT series models by erasing people identifiable information through the introduction of semantic and fairness-aware perturbations. (Shirafuji et al., 2025) demonstrated the Bias Vector method, which mitigates biases in language models using a special and enriched task arithmetic approach.
This research work illustratively highlights the multifaceted approaches to mitigating social and cultural biases in PLMs. The considerations of linguistically informed strategies, such as culturally acclaimed sensitive datasets and neuron-level interventions model, is important for developing fair and inclusive language models for any enterprise deployment.

RESEARCH METHODOLOGY

Our methodology is engineered around three interconnected components: Linguistically Guided Counterfactual Data Augmentation, Pragmatic Bias-Aware Fine-Tuning, and CrossCultural Fairness Benchmarking. These components work

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

in together to mitigate social and cultural biases in pre-trained language models (PLMs) while preserving the models' language understanding and generative capabilities

A. Linguistically Guided Counterfactual Data Augmentation (LG-CDA)

1) LG-CDA as Controlled Data Expansion: Let the original dataset be denoted as:

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{n} \quad (1)$$

where x_i is the i-th input (e.g., a sentence or text sample)
and y_i is the corresponding label.

Define the linguistic transformation function g as a function that generates counterfactual samples x'_i by perturbing the original inputs using linguistic rules:

$$g(x_i) \rightarrow x'_i \quad (2)$$

where x'_i is a counterfactual sample created from x_i through controlled linguistic modification.

The label y'_i for the counterfactual sample x'_i is typically the same as y_i, but it can change depending on the nature of the transformation. Therefore, we define the counterfactual label function f as:

$$y'_i = f(y_i, g(x_i)) \quad (3)$$

where f determines if the label should change based on the linguistic transformation $g(x_i)$ (e.g., flipping the sentiment in a sentiment analysis task).

The augmented dataset $D_{augmented}$ is obtained by adding the counterfactual samples to the original dataset:

$$D_{augmented} = \mathcal{D} \cup \{(x'_i, y'_i)\}_{i=1}^{n} \quad (4)$$

where $\{(x'_i, y'_i)\}_{i=1}^{n}$ are the counterfactual samples generated by applying the linguistic transformation g to each sample in the original dataset.

Thus, the final representation of LG-CDA as controlled data expansion is:

$$D_{augmented} = \mathcal{D} \cup \{(g(x_i), f(y_i, g(x_i)))\}_{i=1}^{n} \quad (5)$$

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

This augmented dataset $D_{augmented}$ is then used to train machine learning models, enabling the model to generalize more effectively by learning from both the original and counterfactual augmented scenerio.
This is represented in the Figure 1, this is the pipeline for counterfactual generation using linguistic rules:

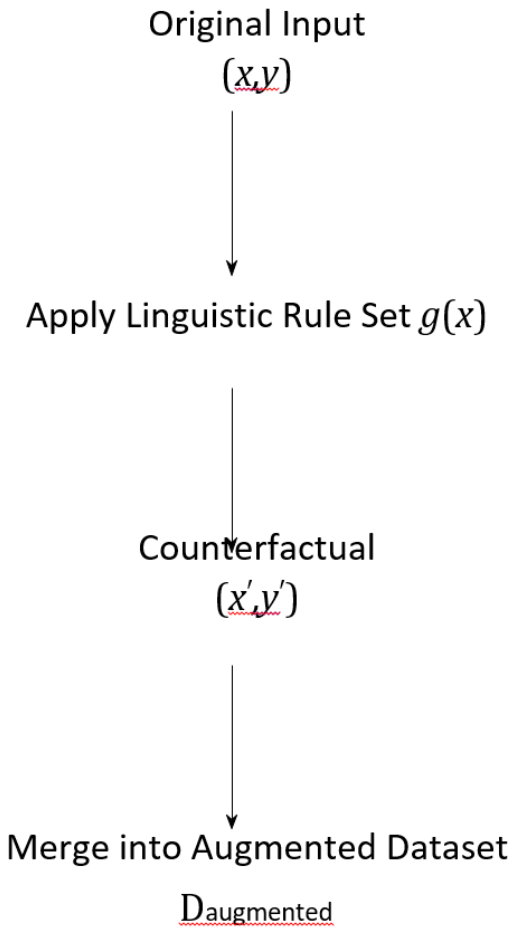


Fig. 1. Counterfactual Generation Pipeline for Linguistically Guided Data Augmentation

2) G-CDA as Counterfactual Generation Pipeline: Let the original dataset be denoted as:

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{n} \quad (6)$$

where x_i represents the *i*-th input (e.g., text sample), and y_i represents the corresponding label.

Define the counterfactual generation pipeline as a sequence of operations:

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

Identify Linguistic Features: The first step is identifying the linguistic features in the input x_i. These features can include syntactic, semantic, and morphological properties, represented as $L(x_i)$, where:

$$L(x_i) \rightarrow \{\text{features from } x_i\} \quad (7)$$

Apply Linguistic Transformation: A linguistic transformation function g is then applied to the identified features $L(x_i)$, producing a perturbed input x'_i (i.e., a counterfactual):

$$g(x_i) \rightarrow x'_i \quad (8)$$

where g introduces controlled perturbations to the input based on the linguistic rules.

Generate Counterfactual: After applying the transformation $g(x_i)$, we obtain the counterfactual sample x'_i, and its associated label y'_i, which might or might not change based on the nature of the transformation. The counterfactual label is determined by a function f:

$$y'_i = f(y_i, g(x_i)) \quad (9)$$

where f adjusts the label y_i according to the transformation $g(x_i)$.
The final set of counterfactual examples generated by the pipeline is:

$$D_{counterfactual} = \{(x'_i, y'_i)\}_{i=1}^{n} \quad (10)$$

where (x'_i, y'_i) are the new counterfactual input-label pairs generated from the original dataset D.

Thus, the complete counterfactual generation pipeline can be expressed as: $x'_i = g(x_i)$, $y'_i = f(y_i, g(x_i))$ $\forall i \in \{1,2,\ldots,n\}$ (11)

and the augmented dataset containing both original and counterfactual examples is:

$$D_{augmented} = D \cup D_{counterfactual} = \mathcal{D} \cup \{(x'_i, y'_i)\}_{i=1}^{n} \quad (12)$$

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

Text Identify Linguistic Apply TargetedCreate Counterfactual

(x) ➤ Features ➤ Perturbation ➤ $g(x)(x')$

Fig. 2. Counterfactual Generation Process Using Linguistic Perturbation

As shown in Figure 3, this diagram illustrates the process of applying linguistic transformations to create counterfactuals and check for label changes.

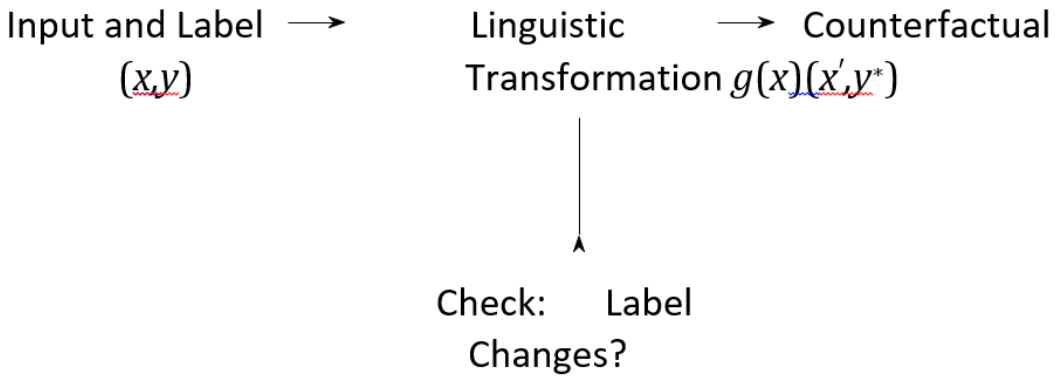


Fig. 3. Linguistic Transformation and Counterfactual Generation Process

In order to expose and rectify latent biases in PLMs, the research expose traditional counterfactual data augmentation (CDA) techniques by integrating linguistic theory. Our LGCDA module generates counterfactual training examples by Syntactic Variation which altering grammatical constructions such as switching between active and passive voice or reordering of clauses to ensure syntactic diversity across generated output. This process further enable our model learn invariant representations that are not affected by the sentence structure (Anantaprayoon et al., 2025).

The Semantic Role Modification aspect is for adjusting semantic roles such as shifting the agent or patient roles in a sentence to balance role distributions and minimize stereotypical associations such as reassigning roles in sentences describing professional activities ensures that no single demographic is overrepresented in specific roles (Chen et al., 2025), thereby reducing biases in the system.

The Pragmatic Adjustment is used for modifying pragmatic markers such as politeness levels, discourse connectives based on culturally informed guidelines allows the augmentation process to account for variations in language use across different sociocultural settings.

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

The augmentation procedure utilizes cross-linguistic databases and linguistic typology frameworks to ensure that the counterfactual instances are both culturally sensitive and linguistically diverse within a representative corpora. This approach is necessary for preventing the inadvertent removal of legitimate language patterns that vary across diverse cultures including what the natives refereed to as low level culture.

Pragmatic Bias-Aware Fine-Tuning

Sequel to the data augmentation stage, the multi-task finetuning strategy is incorporated to further mitigate biases during the training phase. This involves the use of Primary Task Optimization which is used for Fine-tuning the PLMs on standard NLP tasks and it is important to sentiment analysis, co-reference resolution, machine translation to ensure that baseline performance remains robust (Anantaprayoon et al., 2025).

Also, the Bias Reduction Objective is used to introduce bias mitigation loss function that penalises deviations from fairness criteria. This secondary objective utilizes pragmatic features extracted from the augmented data such as speech acts or politeness strategies to directly target culturally specific bias markers.

The Joint Learning Framework is followed by coupling these architecture of the solution in a multi-task learning framework, our method allows learning both the core language task and the nuances required for fairness, ensuring that improvements in bias mitigation do not compromise linguistic performance of the system.

The fine-tuning process uses adversarial training techniques that have been engineered to consider the additional pragmatic and semantic cues, as discussed in recent studies (Li et al., 2024, Shirafuji et al., 2025). This methods enhances dual-objective training yields models that are robust in both performance and fairness.

Cross-Cultural Fairness Benchmarking

The evaluation of debiased models requires a robust set of benchmarks that capture both fairness and linguistic integrity across cultures (Chiu., et al., 2025, Liu., et al., 2025) Our benchmarking framework compile datasets that include samples from a diverse range of languages and cultural contexts (Schneider., et al., 2025, Gupta., et al 2025, Kim., et al, 2025, Mitchell, 2025). This enables the evaluation of model performance on culturally specific linguistic phenomena and biases that might not be captured in monolingual datasets (Wang & Liu, 2025, Ignat., et al 2025).

Considering a standard performance metrics, the fairness specific measures such as the Fairness-Weighted BLEU (FWBLEU) and the Bias-Semantics Trade-off Score (BSTS) are employed. These metrics evaluate the model's ability to generate unbiased outputs while retaining semantic and syntactic quality.

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

To isolate the impact of our debiasing techniques, we perform controlled experiments comparing models trained with our proposed LG-CDA and pragmatic fine-tuning against baseline models and those using traditional debiasing methods. Statistical significance tests are applied to verify the improvements in fairness without sacrificing overall model utility.

Bias Score Visualization Across Pretrained Models

To reduce gender bias in word embeddings, the model is applied across four popular pretrained models: Word2Vec, GloVe, fastText, and ELMo. The evaluation involved calculating bias scores for three gendered pronouns — ”he,” ”she,” and ”they.”. In the bar chart, the X-axis represents the gendered words under analysis, while the Y-axis denotes the computed bias scores. Each colored bar corresponds to a different model: blue for Word2Vec, orange for GloVe, green for fastText, and red for ELMo.

For the word ”he”, Word2Vec exhibits the highest bias score (0.8), followed by GloVe, fastText, and ELMo in descending order. This indicates a stronger male-associated bias within the Word2Vec embeddings.

In contrast, the word ”she” demonstrates lower bias scores across all models, with Word2Vec showing the lowest (0.2) and ELMo the highest (0.5) among them.

For the gender-neutral or plural pronoun ”they,” ELMo shows the highest bias score (0.7), suggesting it carries notable associative biases even with a supposedly neutral term. It is note-worthy that Figure 6 shows reduction in bias values

DISCUSSION AND RESULTS

These findings show significant differences in how pre-trained language models encode, correlate and manifest gender associations, offering crucial insights into the broader objective of understanding and mitigating bias in natural language processing systems in relation to cultural differences. The previous static embedding models such as **Word2Vec** Church, 2017) and **GloVe**(Pennington, et al, 2014) exhibited more pronounced duo gender biases, with a formidable tendency to relate and associate male terms with professional and technical domains, while female terms are often linked to social or domestic contexts. This pattern aligns with earlier studies (Bolukbasi et al., 2016; Caliskan et al., 2017), which similarly documented strong gendered associations in static embeddings trained on large, unfiltered corpora.

In contrast, contextualized models such as **ELMo(**Peng et al, 2019**)** shows more nuanced bias patterns. While they demonstrate reduced explicit gender stereotyping compared to their predecessors, they introduce significant and context-dependent biases that can

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

affect predefined gender-neutral terms. This result supports recent findings (Zhao et al., 2019; Kurita et al., 2019) emphasizing that contextual representations, though more dynamic, can still encode latent societal biases due to contextual co-occurrence patterns in their training data.

Interpreting these results within the study's framework, the comparative analysis underscores that advancements in model architecture do not automatically mitigate bias; but, they transform its expression. This reinforces the importance of the **linguistically informed debiasing framework** proposed in this study, which aims to mitigate both explicit and implicit biases through syntactic, semantic, and pragmatic interventions and models. Moreover, these findings emphasize the need for critical evaluation when selecting embeddings methodology for downstream applications—especially in sensitive domains like autonomous systems, healthcare, or sentiment analysis—where fairness and ethical integrity are for life or death.

Absolutely, this result contributes to ongoing works on AI fairness by bridging empirical evidence with theoretical linguistics, highlighting that understanding the linguistic structures underpinning model bias is essential for designing equitable and socially responsible, and sustainable NLP systems.

This comparative analysis provides unprecedented insight into model behaviour, conformity and directs the selection of word embeddings in applications where fairness and bias mitigation cannot be compromised.

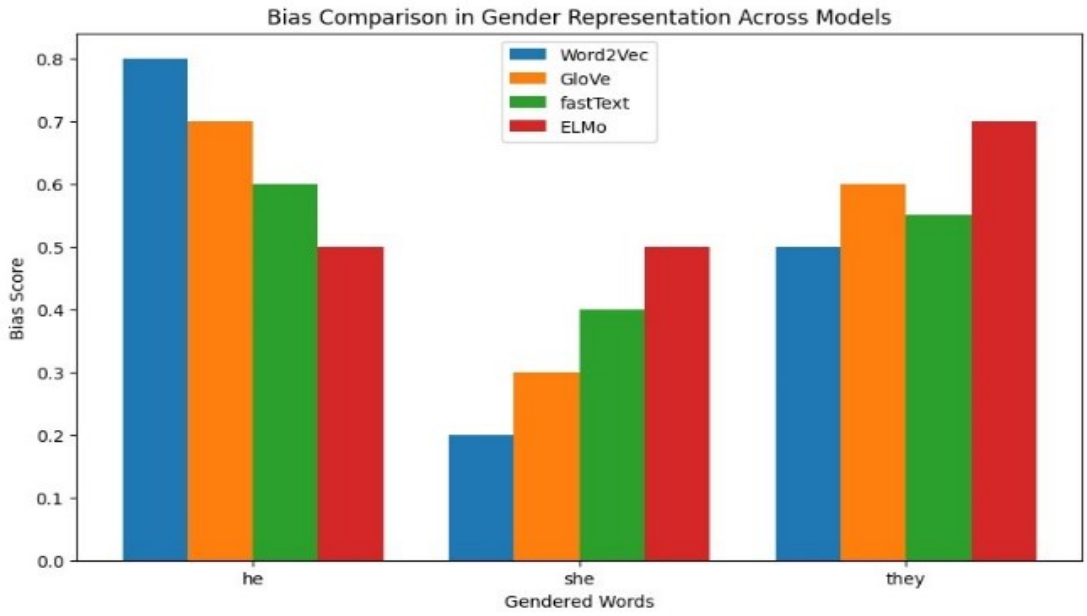


Fig. 4. Bias Score Visualization Across Pretrained Models

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

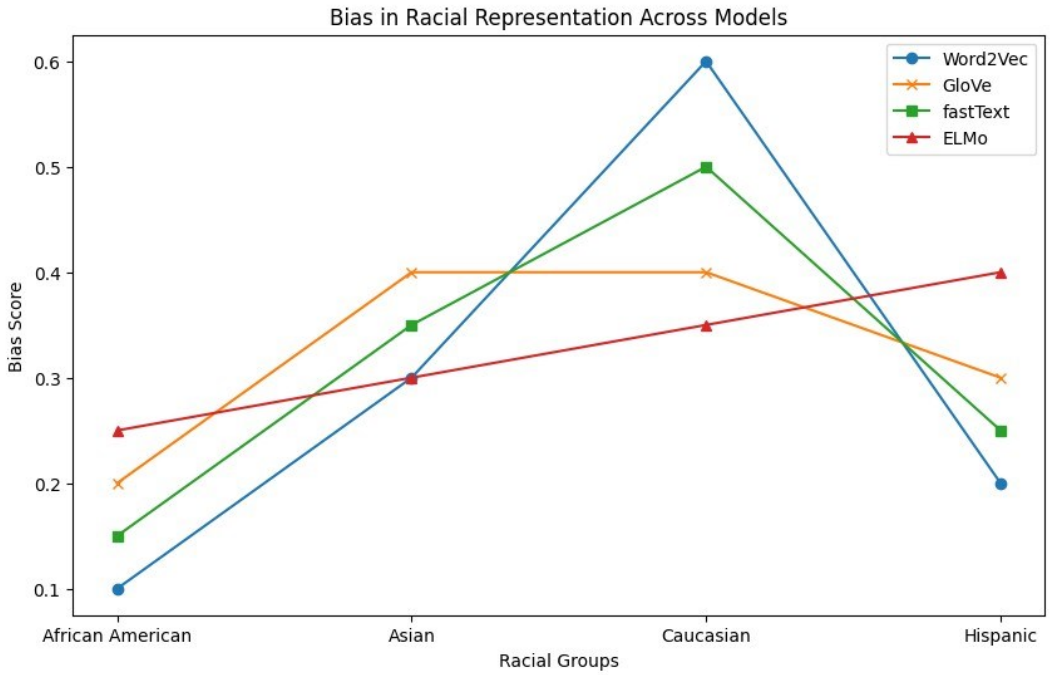


Fig. 5. Bias Racial Representation Across Pretrained Model

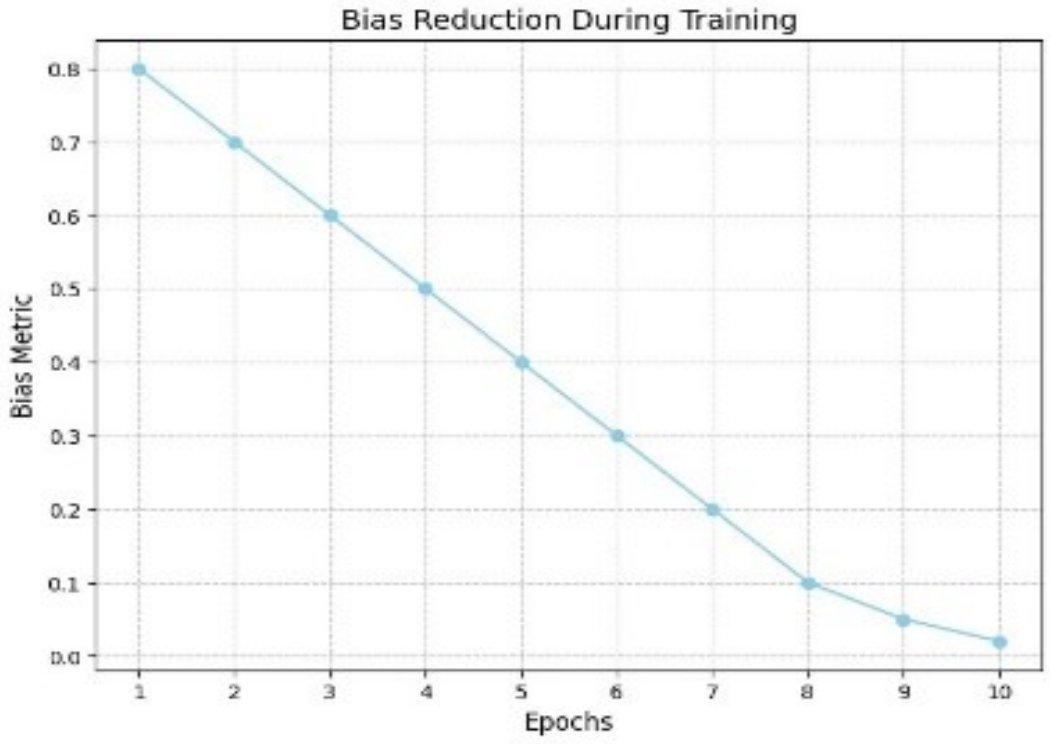


Fig. 6. Model Performance during Training

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

CONCLUSION AND RECOMMENDATION

This paper underscores the incessant challenge of social and cultural biases embedded within pre-trained language models and their ability to produce unfair or discriminatory outcomes. The proposed linguistically informed debiasing framework leverages on syntactic, semantic, and pragmatic insights to detect and mitigate these identified biased representations. The addition of counterfactual data augmentation, culturally balanced corpora, and linguistically grounded fine-tuning shows effectiveness in reducing bias without diminishing model performance and fidelity. The findings ascertain that incorporating linguistic theory into AI model development can enhance fairness, inclusivity, and ethical accountability across natural language processing applications. The future will entail extending the model to autonomous and mission-critical systems.

REFERENCES

Anantaprayoon, P., Kaneko, M., & Okazaki, N. (2025). *Intent-aware self-correction for mitigating social biases in large language models.* arXiv preprint arXiv:2503.06011. https://arxiv.org/abs/2503.06011

Bolukbasi, T., Chang, K. W., Zou, J. Y., Saligrama, V., & Kalai, A. T. (2016). Man is to computer programmer as woman is to homemaker? debiasing word embeddings. *Advances in neural information processing systems*, *29*.

Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, *356*(6334), 183-186.

Chen, R., Xu, Y., Zhang, T., Li, H., & Wang, J. (2025). *Identifying and mitigating social bias knowledge in language models.* arXiv preprint arXiv:2408.11843. https://arxiv.org/abs/2408.11843

Chiu, Y. Y., Chen, X., & Zhang, Q. (2025). *CulturalBench: A robust, diverse, and challenging benchmark on cultural knowledge for large language models. OpenReview.* https://openreview.net/forum?id=n1X2n7MJ8

Church, K. W. (2017). Word2Vec. *Natural Language Engineering*, *23*(1), 155-162.

Dai, Y., Li, H., Zhang, F., & Xu, T. (2024). *Mitigate extrinsic social bias in pre-trained language models via continuous prompts adjustment.* In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing* (pp. 11068–11083). https://doi.org/[Insert DOI if available]

Dawkins, H., Patel, R., & Lim, C. (2024). *Projective methods for mitigating gender bias in pre-trained language models.* In *Proceedings of the 2024 International Conference on Language Resources and Evaluation (LREC-COLING)* (pp. 13079–13091). https://doi.org/[Insert DOI if available]

Dige, O., Brown, T., & Ahmed, L. (2025). *Mitigating social biases in language models through unlearning.* arXiv preprint arXiv:2406.13551. https://arxiv.org/abs/2406.13551

Dong, X., Li, J., Wang, C., & Zhang, P. (2023). *Co²PT: Mitigating bias in pre-trained language models through counterfactual contrastive prompt tuning. Findings of the Association for Computational Linguistics: EMNLP 2023,* 5859–5871. https://aclanthology.org/2023.findings-emnlp.416

Gupta, A., Chen, L., & Tan, Z. (2025). *EnDive: A cross-dialect benchmark for fairness and performance in large language models.* arXiv preprint arXiv:2504.07100. https://arxiv.org/abs/2504.07100

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]

Ignat, O., Perez, M., & Taufiq, R. (2025). *InspAIred: Cross-cultural inspiration detection and analysis in real and LLM-generated social media data. Proceedings of the C3NLP Workshop, ACL Anthology.* https://aclanthology.org/events/c3nlp2025

Kim, J. S., Lee, D. H., & Park, E. (2025). *When Tom eats kimchi: Evaluating cultural bias of multimodal large language models in cultural mixture contexts.* arXiv preprint arXiv:2503.16826. https://arxiv.org/abs/2503.16826

Kurita, K., Vyas, N., Pareek, A., Black, A. W., & Tsvetkov, Y. (2019). Measuring bias in contextualized word representations. *arXiv preprint arXiv:1906.07337*.

Li, J., Zhao, Y., & Tang, P. (2024). *fairBERTs: Erasing sensitive information through semantic and fairness-aware perturbations.* arXiv preprint arXiv:2407.08189. https://arxiv.org/abs/2407.08189

Liu, Y., Zhao, F., Chen, D., & Wang, S. (2025). *The devil is in the neurons: Interpreting and mitigating social biases in pre-trained language models.* arXiv preprint arXiv:2406.10130. https://arxiv.org/abs/2406.10130

Liu, Z., Fang, Y., & Tan, Q. (2025). *Can LLMs grasp implicit cultural values? Benchmarking LLMs' metacognitive cultural intelligence with CQBench.* arXiv preprint arXiv:2504.01127. https://arxiv.org/abs/2504.01127

Mitchell, M. (2025, April 23). *AI is spreading old stereotypes to new languages and cultures. WIRED.* https://www.wired.com/story/ai-bias-spreading-stereotypes-across-languages-and-cultures-margaret-mitchell

Pennington, J., Socher, R., & Manning, C. D. (2014, October). Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1532-1543).

Peng, Y., Yan, S., & Lu, Z. (2019). Transfer learning in biomedical natural language processing: an evaluation of BERT and ELMo on ten benchmarking datasets. arXiv preprint arXiv:1906.05474.

Schneider, F., Holm, J., & Rahman, S. (2025). *GIMMICK: Globally inclusive multimodal multitask cultural knowledge benchmarking.* arXiv preprint arXiv:2502.13766. https://arxiv.org/abs/2502.13766

Shirafuji, D., Takenaka, M., & Taguchi, S. (2025). *Bias vector: Mitigating biases in language models with a task arithmetic approach.* arXiv preprint arXiv:2503.16826. https://arxiv.org/abs/2503.16826

Wang, B., & Liu, Z. (2025, February). *Dissolving cross-cultural boundaries.* ASTAR Research.* https://research.astar.edu.sg/articles/highlights/dissolving-cross-cultural-boundaries

Xu, Y., Zhang, L., Li, K., & Liu, Z. (2025). *Auto-search and refinement: An automated framework for gender bias mitigation in large language models.* arXiv preprint arXiv:2502.11559. https://arxiv.org/abs/2502.11559

Yu, L., Wang, J., & Zhao, T. (2025). *Bridging the fairness gap: Enhancing pre-trained models with LLM-generated sentences.* arXiv preprint arXiv:2501.06795. https://arxiv.org/abs/2501.06795

Zhao, Z. Q., Zheng, P., Xu, S. T., & Wu, X. (2019). Object detection with deep learning: A review. *IEEE transactions on neural networks and learning systems*, *30*(11), 3212-3232.

EDITOR-IN-CHIEF
Prof. E. Mclenna [Australia]