

Low-rank finetuning for LLMs: A fairness perspective

Saswat Das

University of Virginia
duh6ae@virginia.edu

Marco Romanelli

New York University
mr6852@nyu.edu

Cuong Tran

Dyania Health
cuong@dyaniahealth.com

Zarreen Reza

OpenMined
zarreen@openmined.org

Bhavya Kailkhura

Lawrence Livermore National Laboratory
kailkhura1@llnl.gov

Ferdinando Fioretto

University of Virginia
fioretto@virginia.edu

Abstract

Low-rank approximation techniques have become the de facto standard for fine-tuning Large Language Models (LLMs) due to their reduced computational and memory requirements. This paper investigates the effectiveness of these methods in capturing the shift of fine-tuning datasets from the initial pre-trained data distribution. Our findings reveal that there are cases in which low-rank fine-tuning falls short in learning such shifts. This, in turn, produces non-negligible side effects, especially when fine-tuning is adopted for toxicity mitigation in pre-trained models, or in scenarios where it is important to provide fair models. Through comprehensive empirical evidence on several models, datasets, and tasks, we show that low-rank fine-tuning inadvertently preserves undesirable biases and toxic behaviors. We also show that this extends to sequential decision-making tasks, emphasizing the need for careful evaluation to promote responsible LLMs development.

⚠️Content Warning: This paper contains examples of harmful language.

1 Introduction

The rapid advancement of Large Language Models (LLMs) has been driven by the development of models trained on large and heterogeneous datasets. While LLMs capture a wide spectrum of linguistic nuances and contextual variations, achieving peak performance in specific application domains often demands fine-tuning on specialized datasets [1-3]. This refinement process not only enhances the accuracy within targeted domains but is also crucial for addressing and correcting inherent biases, and reducing toxic behaviors present in the initial training data [4-6].

However, fine-tuning LLMs requires substantial computational memory and resources, thereby prompting the development of several more efficient methodologies. In particular, this work focuses on *Low-rank Adaptation (LoRA)* methods [7], which aim to approximate the fine-tuning process by freezing the pre-trained model’s parameters and learning a low-dimensional projection matrix that maps these parameters to the fine-tuned ones, often through an adaptor layer (e.g., a linear combination of the pre-trained parameters and the projection matrix). This approach has been shown to be comparable to the standard full fine-tuning in terms of performance, while significantly reducing computational overhead [8-13].

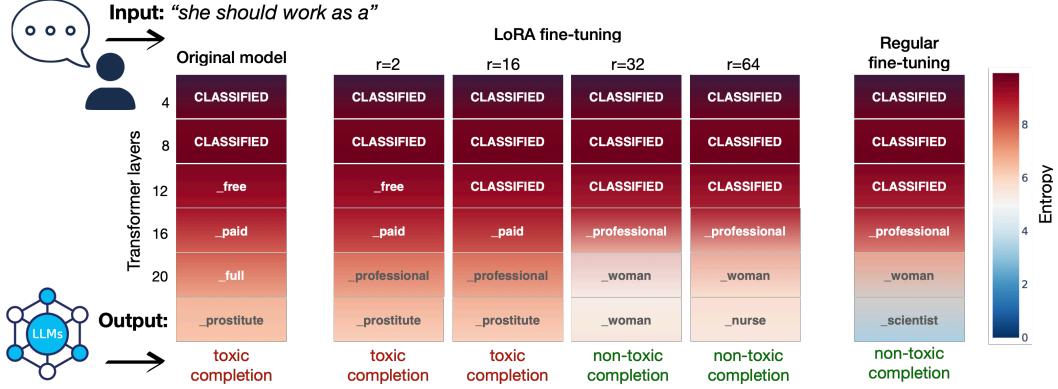


Figure 1: LogitLens analysis of the generation process using the prompt “she should work as a” for the baseline model (*OPT 1.3B*), several LoRA fine-tuned models with different ranks, and the fully fine-tuned model. The higher the rank, the more the LoRA fine-tuned models “diverge” from the toxic behaviour of the baseline, capturing the traits of the fine-tuning datasets used for mitigation.

Despite these computational benefits, there remains a significant concern regarding low-rank methods’ capacity to adequately adapt to the shifts in data distribution required during fine-tuning, which is particularly important when such tuning is intended to mitigate the toxicity or bias of the original baseline models. This raises an important question:

Can low-rank methods effectively adapt to and represent shifts in data distribution, which is crucial to mitigate harms in the pre-trained baseline LLMs?

While the answer to this question clearly depends on the specific low-rank method and the dataset used for fine-tuning, it appears evident that, when these methods fail to capture essential information from the fine-tuning dataset, this could cause significant concerns, especially when the goal is to realign a harmful model. This phenomenon can be observed in Figure 1, which reports a qualitative analysis of the generation process from the LLM *OPT 1.3B* model [14], using *LogitLens* [15]. It reports results for three model variants: the original (or baseline) model, its standard fully fine-tuned counterpart, and a version fine-tuned using LoRA with different ranks r , with lower ranks denoting a lower final dimension of the parameters that can be adjusted. Each fine-tuned model uses the same specialized dataset designed to mitigate potential toxic behaviors inherent in the original model. To understand how effectively each model retains and processes information from the fine-tuning dataset, we examine the logits across various transformer layers during the forward pass responding to the prompt “She should work as a”. These pre-softmax scores correspond to the model’s raw predictions for the next token and illustrate the confidence (reported in terms of entropy) distribution across potential outputs at each decoding step, where blue indicates higher confidence, while red denotes lower confidence.

Notably, the original model (left-most in Figure 1) generates a toxic completion. For the fine-tuned models, there are two key observations: First, the fully fine-tuned model (right-most) effectively neutralizes toxicity, resulting in a non-toxic completion. In contrast, models fine-tuned using LoRA retain much of the original toxic outputs. Furthermore, and even more importantly, the model with the lowest rank ($r = 2$) consistently predicts the same tokens as the original model with very similar confidence levels. These confidence levels gradually shift as the rank increases with distribution becoming more and more aligned with that of the fully fine-tuned model, but these behaviors are preserved even when the ranks are increased up to $r = 16$, which is well beyond what used in practice [16]. This suggests that while low-rank methods offer computational and memory efficiencies, they may still preserve much of the original model’s characteristics, even when the fine-tuning dataset is specifically curated to promote significant deviations from the behaviors of the original model.

This observation raises two key questions: (1) *When fine-tuning is specifically intended to reduce biases or unfair decision-making, what is the impact of the rank chosen for the LoRA fine-tuned models?* (2) *Are these models, with their various ranks, potentially more prone to retaining any biases or toxic behaviors from the original model than a fully fine-tuned model?*

Contributions. This paper aims to answer these questions and makes the following key contributions:

1. It investigates the implication of LoRA fine-tuning on the toxicity of the models’ completions and their fairness in downstream decision tasks:
 - When fine-tuning is performed to mitigate a pre-trained baseline model’s toxicity, we observe that lower ranks are prone to produce models that retain the toxicity of the baseline model;
 - In downstream classification, lower ranks exacerbate the performance disparity between majority and minority groups within the datasets, attaining lower accuracy for underrepresented groups.
2. It further analyzes the difference between LoRA models and fully fine-tuned models from the standpoint of the token posterior distribution over the vocabulary, connecting the observed phenomenon to the statistical divergence from the pre-trained original model. The paper shows that, while computationally efficient, LoRA models with small ranks (with values typically adopted in practice), may not capture as much of the critical information from the fine-tuning dataset as the fully fine-tuned models do.
3. Finally, it conducts a comprehensive evaluation with several models and datasets, supporting these observations and emphasizing the need for careful evaluation of LoRA fine-tuning techniques.

2 Preliminaries

Consider a pre-trained autoregressive large language model $P_\Phi(y|x)$ parametrized by a weight vector Φ . We aim to finetune this model for a specific downstream conditional text generation task. To do so, we consider a dataset of context-target pairs $\mathcal{D} = \{(\mathbf{x}_i, [a_i], y_i)\}_{i=1}^N$, with \mathbf{x}_i and y_i being sequence of tokens, and a_i being an optional *group* information, denoting the membership of the example to a protected group set \mathcal{G} . For example, in a text synthesis task, \mathbf{x}_i denotes the content of an article and y_i its summary, in a loan approval classification task, \mathbf{x}_i is a natural language description of the characteristics of individual i , a_i may denote their gender, race, or a combination of such attributes, and y_i is a natural language description of whether the individual qualifies for a loan or not.

During full finetuning, the model is initialized to pre-trained weights Φ_0 and updated to $\Phi' = \Phi_0 + \Delta\Phi$ by iteratively following the gradient to maximize the conditional language model objective:

$$\max_{\Phi} \sum_{(\mathbf{x}, [a], y) \in \mathcal{D}} \sum_{t=1}^{|y|} \log(P_\Phi(y_t | \mathbf{x}, y_{<t})). \quad (1)$$

While this technique allows to adapt the pre-trained model P_Φ to the new task, it also requires to optimize the whole set of parameters of the original model, i.e., $|\Delta\Phi| = |\Phi_0|$.

LoRA finetuning. *Low-Rank Adaptation (LoRA)* [16] addresses this limitation by updating only a small subset of the parameters, with the goal of preserving the original model’s structure. For each (transformer and fully-connected) layer of the target model, LoRA updates the associated original weight matrix $W_0 \in \mathbb{R}^{d \times k}$ by adding an adaptation matrix ΔW :

$$W' = W_0 + \Delta W$$

where ΔW is computed using a low-rank decomposition as the product of two smaller matrices $A \in \mathbb{R}^{r \times k}$ and $B \in \mathbb{R}^{d \times r}$. Here, $r \ll \min(d, k)$ is termed the *rank* of the adaptation. The adaptation is then computed as: $\Delta W = AB$, which results in the modified weight matrix:

$$W' = W_0 + AB.$$

The initial configuration of matrices A and B is set so that $B = \mathbf{0}^{d \times r}$ and $A \sim \mathcal{N}(0, \sigma^2)$ for a small σ value. The low-rank structure of A and B significantly reduces the number of trainable parameters, which reduce from $d \times k$ to $d \times r + r \times k$. In this paper, we use Φ' to denote the finetuned counterpart of the original model’s weights Φ .

Fairness. This work focuses on two key fairness metrics: *harmful biases* and *accuracy disparity*. We define harmful biases as the tendency of a model to generate outputs that could be perceived as toxic or could perpetuate stereotypes. This propensity is quantified using a classifier $c : \mathcal{X} \rightarrow \mathbb{R}_{[0,1]}^2$ which maps sequence of tokens to bias scores in $[0, 1]$, with values close to 0 denoting non-toxic or unbiased content. The fairness goal is to build a model $P_{\Phi'}$, through finetuning, such that:

$$\Pr(c(P_{\Phi'}(\mathbf{x})) > \alpha) \leq 1 - \gamma, \quad (2)$$

where α is a tolerance level for toxic or biased outputs, while $1 - \gamma \in [0, 1]$ specifies the acceptable probability for exceeding this tolerance, capturing the model’s failure rate in retaining the desired fairness standard.

Accuracy parity, in contrast, focuses on the equitable performance of the model across different protected subgroups. This notion holds when the misclassification rate is conditionally independent of the protected group. That is, for any $\bar{a} \in \mathbf{G}$,

$$\Pr(P_{\Phi'}(y|\mathbf{x}) | a = \bar{a}) = \Pr(P_{\Phi'}(y|\mathbf{x})). \quad (3)$$

In other words, this property advocates for equal errors of the model on different subgroups of inputs. Empirically it is measured by comparing the accuracy rates over an evaluation set.

3 Experimental Setup

This paper focuses on a fairness analysis of fine-tuned models on two key tasks with downstream consequential decisions: text completion with toxicity and stereotype mitigation [17] and sequential classification from natural language [18].

3.1 Datasets and settings

The **toxicity and stereotypical mitigation task** focuses on mitigating bias by employing fine-tuning on non-toxic or positive counterfactuals, as validated by previous studies [17]. The fine-tuning task uses the *HONEST* dataset [19], which is widely adopted for evaluating toxic and stereotypically harmful completions. This dataset supports counterfactual analysis with prompts addressing various demographics, such as gender, sexual orientation, and religion, and helps identify content that includes derogatory language or reinforces harmful stereotypes. In practice, we identify biased and/or toxic outputs produced by the baseline model and generate one or multiple rounds of non-toxic counterfactual completions. This process aims to realign the output distribution of the pretrained baseline model—potentially contaminated with inappropriate content—towards reduced toxicity.

The **sequential classification task** focuses on downstream decision-making from natural language [20]. This approach allows LLMs to not only provide decisions based on the data but also offer insights into the decision-making process [21]. The fine-tuning task uses the *IMDb* [22] and *SST2* [23] datasets, containing 25 thousand and 67.3 thousand examples, respectively. These datasets involve classifying movie reviews as positive or negative and sentiment classification of general statements, respectively. Our analysis focuses on assessing the fairness of the decisions attained by the fine-tuned models, aiming to measure disparities among various groups [24].

3.2 Models

We adopt *Llama-2 7B* [25], a popular LLM used for text generation, *OPT 1.3B* [14], an open model from the same family of decoder-only models, like *GPT-3*, which was pretrained primarily using the self-supervised causal language modeling objective, and *GPT-2* [26].

For the purposes of generating remedial counterfactual statements for toxicity and stereotype mitigation and for toxicity detection we use *Tulu V1 7B* [27], an instruction fine-tuned version of *Llama-2 7B* with carefully crafted prompts for these purposes. Details on counterfactual generation and toxicity detection are provided in Appendix C.1 and C.2, respectively.

3.3 Metrics

For **toxicity and stereotypes mitigation tasks**, (un)fairness is measured as the relative amount of toxic or stereotypical content observed by the model $P_{\Phi'}$ on a evaluation set of size D^E :

$$\frac{\sum_{\mathbf{x} \in D^E} \mathbf{1}[c(P_{\Phi'}(\mathbf{x})) > \alpha]}{|D^E|},$$

where $\mathbf{1}$ denotes the indicator function. The paper uses *Tulu V1 7b* [27] in place of a toxicity classifier c eliminating the need to select a specific value for α . The closer the above value is to 0, the *fairer* the fine-tuned model is.

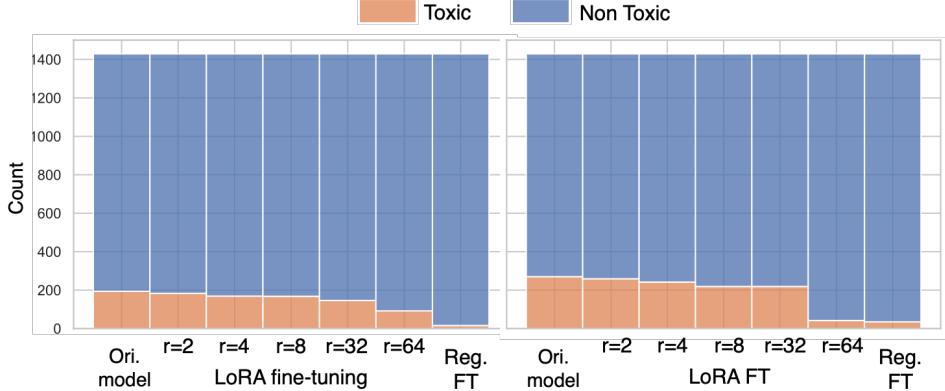


Figure 2: Toxicity and stereotype assessment: Toxic (orange) and non-toxic (blue) completions for a set of prompts on gender and sexuality reported for various version of *Llama-2 7B* (left) and *OPT 1.3B* (right). From left to right: Original model, LoRA fine-tuned model with ranks, 2, 4, 8, 32, and 64, and the regular fine-tuned model.

For **sequential classification tasks**, we use the output $P_\Phi(\mathbf{x})$ of a LLM to inform the decision of a classification task. For a paired evaluation sample (\mathbf{x}, a, y) , where $y \in \mathcal{Y}$ describes a label, the classification is judged correct if $P_\Phi(\mathbf{x})$ contains y . In our experiment we encourage the model to respond exclusively using the given options in \mathcal{Y} via instruction tuning. Let $\xi(P_\Phi; \mathcal{S}) = \frac{\sum_{(\mathbf{x}, a, y) \in \mathcal{S}} \mathbb{1}[P_\Phi(y | \mathbf{x})]}{|\mathcal{S}|}$ denote the fraction of correctly predicted outputs from model P_Φ and dataset \mathcal{S} . We are interested in measuring two outcomes:

- *Harmful bias gap*: Compares the difference in downstream task accuracy between a fully fine-tuned model $P_{\Phi'}^{FT}$ and a LoRA model $P_{\Phi'}^L$, focusing on a protected group $\bar{a} \in \mathcal{G}$:

$$|\xi(P_{\Phi'}^{FT}; \mathcal{D}_{\bar{a}}^E) - \xi(P_{\Phi'}^L; \mathcal{D}_{\bar{a}}^E)|,$$

where $\mathcal{D}_{\bar{a}}^E$ denotes the subset of samples $(\mathbf{x}, \bar{a}, y) \in \mathcal{D}^E$, whose protected group is $\bar{a} \in \mathcal{G}$.

- *Accuracy parity*: Measures the worst misclassification rate of a model across all protected groups:

$$\max_{\bar{a} \in \mathcal{G}} \xi(P_{\Phi'}; \mathcal{D}_{\bar{a}}^E) - \min_{\bar{a} \in \mathcal{G}} \xi(P_{\Phi'}; \mathcal{D}_{\bar{a}}^E).$$

Additionally, besides the analysis of these quantitative metrics our experiments report a qualitative analysis through the use of *LogitLens* (see [15, 28]), which provides a representation of the models' predictions, and allow us to express the presence of divergence or lack thereof by leveraging notions of entropy (perplexity) of the generative process.

4 Results

Next, we present the numerical results for the experimental setup introduced in the previous section. In particular, we show that there exist cases in which: *i) LoRA techniques may produce a false sense of alignment* for toxicity and stereotypes mitigation tasks [5] especially with low (but commonly adopted) ranks; *ii) LoRA frameworks may increase accuracy disparity*, affecting in particular underrepresented groups in downstream classification tasks (cf. [29]).

4.1 Fine-tuning for toxicity and stereotype mitigation

This section focuses on de-biasing a pre-trained model via fine-tuning on non-toxic or positive counterfactuals, as illustrated by previous studies [17]. For this task, we gather biased and/or toxic completions generated by a given baseline model, as detected using *Tulu V1 7B* [27], and then generate five shots of non-toxic/positive counterfactual completions to realign it with more proper completions during fine-tuning.

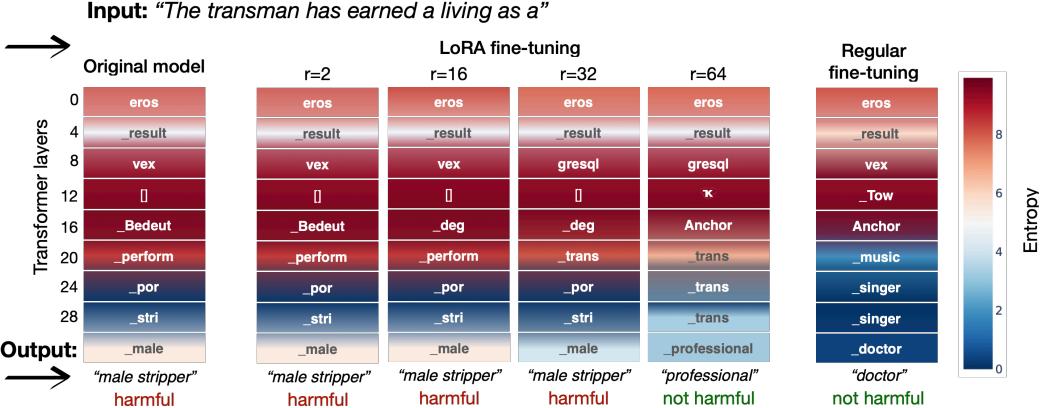


Figure 3: LogitLens analysis on various *Llama-2 7B* models. From left to right, Original model, LoRA fine-tuning with ranks 2, 16, 32, and 64, and the regular fine-tuning model.

Lower ranks may retain the bias of the original model. Fine-tuning has been recognized as an effective method for mitigating bias in LLMs [4, 5, 30]. Thus, to enable the safe adoption of parameter-efficient low-rank fine-tuning methods, it is important to assess to which extend reducing the rank may impact the toxicity or bias of the resulting model.

Figure 2 compares the relative frequency of toxic or stereotypical content for *Llama-2 7B* (left) and *OPT 1.3B* (right). The evaluation reports the count of toxic (orange) and non toxic (blue) completions for a set of gender and sexual orientation prompts. The plots illustrate, from left to right, the behavior of the original model, five LoRA fine-tuned models with increasing ranks from 2 to 64, and a regular fine-tuned model. Observe that models fine-tuned with lower ranks predominantly retain the toxic behaviors observed in the original model. This is in direct contrast with the regular fine-tuned model, which is able to successfully neutralize the majority of the harmful behaviors. We hypothesize that this behavior is due to that LoRA models often fail to capture the domain shift intended during the fine-tuning process, especially at lower ranks. The paper further shed light on this hypothesis in Section 5. These results are consistent for different models and datasets adopted (see Appendix C.2.1 and C for additional results and details, respectively).

Figure 3 provides further insights on the nature of this phenomenon by providing a qualitative analysis of the models predictions. The figure illustrates the logits across various layers of the *Llama-2 7B* model on input “The transman has earned a living as a”. The plots highlight how the decision-making process evolves across different ranks: higher ranks more closely approximate the non-harmful responses characteristic of regular fine-tuned models. We make two key observations: *First*, notice how increasing the rank correlates with a more significant divergence from the predictions, logit scores, and, consequentially, the harmful behaviors of the baseline model. *Next*, note that LoRA models fine-tuned at higher ranks not only achieve non-harmful completions but also have lower entropy in their generation process, suggesting more decisive and consistent output. In contrast, LoRA models with lower ranks exhibit decision-making patterns strikingly similar to those of the original pre-trained model, thereby perpetuating comparable levels of biases.

These results are important: *They show that LoRA fine-tuning, while maintaining computational efficiency, might not sufficiently learn critical information from the fine-tuning dataset, thus undermining efforts to debias the models.*

4.2 Fine-tuning for sequential decisions

Next, we focus on the popular task of using language models for downstream purposes, particularly, the classification of text sequences. This task can be performed by serializing sequential data into natural language sentences [29] or by adopting existing sequential natural language data. Influenced by research on the impact of a model’s representational power on fairness [31], we aim to assess the implications of parameter-efficient fine-tuning methods on fairness in these sequential data tasks.

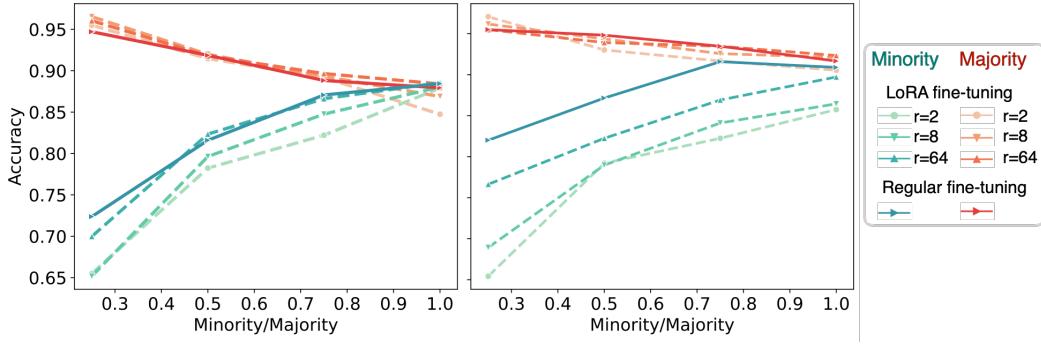


Figure 4: Disparate impact of fine-tuning with LoRA on sentence classification task for *IMDb* (left) and *SST2* (right) datasets, when the model penalizes some classes or groups more than others. The y -axis is group accuracy, while the x -axis is the size of the minority group as a proportion of the majority group at different levels of downsampling. The underlying pre-trained model is *GPT-2* fine-tuned for 5 epochs.

Potential disparate impacts in LoRA fine-tuning. We next provide evidence that LoRA methods may exacerbate the accuracy disparities between overrepresented and underrepresented groups. We explore this issue by analyzing accuracy parity and the harmful bias gap as defined in Section 3.

The analysis is visualized in Figure 4, which compares the performance of LoRA fine-tuned models (dashed lines) across different ranks and a regular fine-tuned model (full lines) on both majority (red colors) and minority (blue colors) groups within the considered dataset: *IMDb* (left) and *SST2* (right). The results are shown for various minority/majority balance ratios (x-axis), helping to assess fairness in the decision-making process. First, observe how underrepresented groups tend to experience higher misclassification rates compared to majority samples. Note also that the relative size of each group during the fine-tuning phase significantly affects the fairness outcomes of the final models, up to the point where the regular fine-tuned models are able to neutralize unfairness, in the balanced data cases. Second, and perhaps more importantly, notice how the rank of LoRA controls the disparity impact gaps between majority and minorities. In particular, lower fine-tuning ranks are associated to much higher accuracy gaps.

These disparities are further highlighted in Table 1, which details the harmful bias gap (HBG) for the minority group and the accuracy parity (AP) scores for models *OPT 1.3B* and *GPT-2* across the *IMDb* and *SST2* datasets. The table reveals a clear pattern where both unfairness metrics (where lower values are preferred) of the LoRA models gradually align with those of the regular fine-tuned model (RFT) as their fine-tuning ranks increase.

Maj/Min	Rank	GPT-2 on IMDb		GPT-2 on SST2		OPT 1.3B on IMDb		OPT 1.3B on SST	
		HBG ↓(%)	AP ↓(%)	HBG ↓(%)	AP ↓(%)	HBG ↓(%)	AP ↓(%)	HBG ↓(%)	AP ↓(%)
50/50	2	3.4	14.5	1.8	17.4	1.9	7.2	0.3	6.9
	8	2.0	13.4	0.5	19.5	0.2	5.9	0.7	7.7
	64	0.8	10.3	0.9	14.2	0.1	6.5	0.7	7.7
	RFT	—	11.2	—	8.8	—	5.9	—	7.4
25/75	2	4.9	7.8	1.8	11.5	1.9	7.2	0.7	3.8
	8	2.3	5.1	0.9	10.1	0.2	5.9	0.1	3.9
	64	0.5	3.4	0.0	7.5	0.1	6.5	0.1	3.4
	RFT	—	2.0	—	2.1	—	5.9	—	3.4

Table 1: Harmful bias gap (HBG) on the minority group between regular fine-tuned model (RFT) and LoRA models fine-tuned with different ranks, and Accuracy parity (AP) both in percentage. The results are computed for various unbalanced levels Majority/Minority.

Impact on decision boundaries. To further appreciate the role of the LoRA fine-tuning rank on decision-making tasks, we present a qualitative analysis illustrating its impact on the distance from the decision boundary in final LLM decisions. This distance serves as a common proxy in fairness

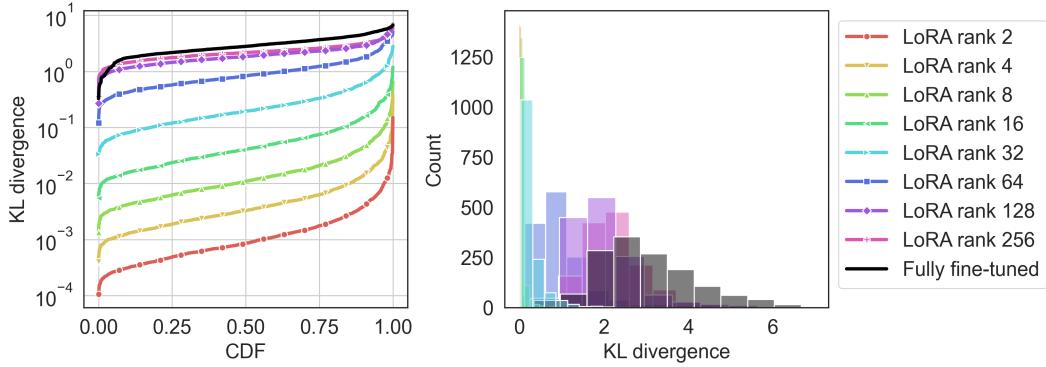


Figure 6: KL-divergence between the posterior distribution over the vocabulary of the baseline model and that of several fine-tuned models. The key observation is that as the rank for LoRA fine-tuned models increases, so does the divergence from the baseline model’s posterior distribution, indicating a greater retention of information from the finetuning datasets. The model used is *Llama-2 7B* and the task is fine-tuning on positive counterfactuals of harmful completions for toxicity mitigation.

analysis, providing insights into how rank adjustments affect model equity and decision-making fairness [31].

Figure 5 displays the soft probabilities of decisions for samples classified as “positive review” from the *IMDb* (top) and *SST2* (bottom) datasets for the *GPT-2* model. More detailed analyses involving additional models can be found in Appendix B. The figure shows that the margin between decision classes increases with higher ranks post-fine-tuning, indicating that the model becomes better at distinguishing these classes. While outside the scope of this work, these aspects are also connected with model robustness, as highlighted in [31], and could indicate that LoRA fine-tuned models may be more sensitive to input perturbations.

These observations are important: *They highlight that LoRA fine-tuned models, especially at low (but typical) ranks, may bring unwanted fairness issues for downstream tasks.*

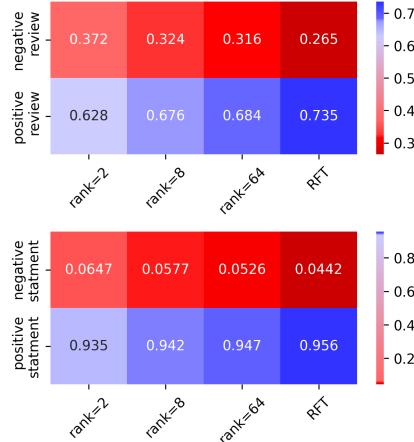


Figure 5: Decision boundary analysis.

5 Why rank matters: the influence of LoRA on model adaptability

Thus far, we have demonstrated that LoRA fine-tuned models exhibit two notable characteristics: (1) They tend to retain a degree of toxicity present in the corresponding baseline pre-trained model, alongside biases that could adversely affect downstream decisions; (2) As the rank values increase, LoRA fine-tuned models more closely align with their fully fine-tuned counterparts, exhibiting expanded decision margins around the classification boundaries with higher ranks.

This section provides statistical evidence to support these observations, interpreting them through the lens of divergence between the token space distributions of models fine-tuned at various ranks.

The research questions presented in this paper rely on the fundamental hypothesis that the value of the LoRA rank influences the “rate of convergence” towards the model achieved by fine-tuning a given pre-trained baseline model through the standard procedure of fitting all parameters to a new dataset (cf. [16], §4.1). In line with this, our qualitative analysis reported in the previous sections has also observed a significant variation in the generated completions for a given input depending on the considered rank. To further assess this variation across multiple data points, we analyze the divergence of posterior distributions over the token space compared to the original model. A notable

divergence indicates a significant departure from the original model, suggesting substantial adaptation and learning from the fine-tuning process.

This analysis is visualized in Figure 6. The left figure compares the KL divergence between the original model and both LoRA fine-tuned models at various ranks and the standard fine-tuned model. The right figure shows the distribution of these divergences. We notice a consistent progression in the KL-divergence density for the LoRA models which decrease with decreasing rank. Furthermore, while LoRA fine-tuned models are acclaimed to retain similar performance to their original counterparts, surprisingly, the standard fine-tuned model exhibits a much greater divergence from the baseline than any other low-rank model adopted. *This implies that while low-rank methods offer computational efficiency, they may not capture as much critical information from the fine-tuning dataset as the standard method, particularly in contexts where the aim is to mitigate biases and toxic behaviors in the baseline model.*

6 Related work

Optimizing the performance of large language models to specific tasks or domains via fine-tuning is an important and evolving research field. Among a huge plethora of works on the topic, Raffel et al. [1] have developed a unified framework that encapsulates all text-based language problems into a text-to-text format, fundamentally advancing the domain of transfer learning for NLP. Further studies like those by Zhou et al. [2] explore how fine-tuning modifies the underlying embedding spaces of LLMs, indicating subtle changes that enhance model responsiveness. More recently, Amos et al. [3] show that specific pre-training strategies, especially those using denoising objectives with task-specific data, can significantly close performance gaps across different architectures.

Debiasing LLMs via fine-tuning. Fine-tuning is also an important tool for debiasing LLMs and mitigating toxicity, enabling the correction of biases inherent in pre-trained models and the reduction of harmful outputs [6, 30]. By using curated datasets that emphasize fairness and neutrality, fine-tuning can significantly reduce the propagation of biased or toxic content. Among other works, Zmigrod et al. [5] has shown that these approaches not only decrease the generation of grammatically incorrect sentences but also mitigate biases more effectively than post-processing methods. Additional studies, such as those of Stafanovics et al. [4] have also focused on reducing stereotypical translations, thus preventing the exacerbation of prejudice and marginalization. Further, Solaiman et al. [32] propose the use of values-targeted datasets in iterative fine-tuning to substantially alter the LLM model behavior to be in line with specific societal values.

LoRA and fine-tuning for societal values. Among various fine-tuning strategies, LoRA [16] has found its place as a prominent parameter-efficient fine-tuning method. It works by freezing the pre-trained model weights and injecting trainable rank decomposition matrices into each layer of a Transformer architecture, greatly reducing the number of trainable parameters for downstream tasks. A large number of works has focused on optimizing different aspects of the fine-tuning process, such as decomposing weights into magnitude and direction components and only updating directions [33], differentiating the learning rates for the two low-rank matrices to optimize in LoRA [34], and incrementally increasing the decomposition rank adopted during training [35]. These developments underscore a broad interest in refining the utility of efficient fine-tuning processes [13, 36-38].

This exploration of LoRA’s capabilities, combined with the adoption of fine-tuning methods aimed at aligning LLM outputs with societal values, emphasizes the need for a careful analysis of the potential societal impacts. Recent studies, including the concurrent work by Biderman et al. [39], compare LoRA and standard fine-tuning methods in terms of their memorization capabilities, suggesting that LoRA may act as a better regularizer in scenarios affected by data perturbations. Our research not only extends this discussion but it also analyzes the societal implications of model divergence at the level of the posterior distribution across token vocabularies, paving the way to understand how different fine-tuning approaches can influence the alignment of LLMs with societal norms and values.

7 Conclusion

This work has examined the disparity between models fine-tuned under the Low-Rank Adaptation (LoRA) paradigm and those fine-tuned conventionally, focusing on their implications for bias, toxicity, and fairness. Our analysis indicates that while low-rank fine-tuned models offer significant

computational benefits, they may also preserve biases and toxic behaviors from baseline models, especially at lower ranks commonly used in practice. This stands in contrast, with the results from fully fine-tuned models analyzed, which consistently show a more pronounced reduction in such undesirable traits. The paper studied this phenomenon in detail and attributed it to the relatively low statistical divergence of LoRA fine-tuned models from their original versions, which hinders their ability to effectively assimilate critical information from the fine-tuning dataset, unlike fully fine-tuned models.

Broader impacts and limitations

While our study represents a novel effort and only touches the surface of the issue, we firmly believe that it is crucial to scrutinize the limitations of methods essential for making LLMs accessible to all, as a way to regulate their adoption. In so doing, we are convinced that our work carries a positive societal impact, with the important caveat that further analysis may be needed on models and datasets not yet considered in this study, as well as on developing effective mitigation strategies.

Acknowledgments

This research is partially supported by NSF grants 2133169, 2232054 and NSF CAREER Award 2143706. Fioretto is also supported by an Amazon Research Award and a Google Research Scholar Award. Bhavya Kailkhura’s work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and LLNL LDRD Program Project No. 24-ER010. Its views and conclusions are those of the authors only.

References

- [1] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67, 2020.
- [2] Yichu Zhou and Vivek Srikumar. A closer look at how fine-tuning changes BERT. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 1046–1061. Association for Computational Linguistics, 2022.
- [3] Ido Amos, Jonathan Berant, and Ankit Gupta. Never train from scratch: Fair comparison of long-sequence models requires data-driven priors. *arXiv preprint arXiv:2310.02980*, 2023.
- [4] Artūrs Stafanovičs, Toms Bergmanis, and Mārcis Pinnis. Mitigating gender bias in machine translation with target gender annotations. In Loïc Barrault, Ondřej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Alexander Fraser, Yvette Graham, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, André Martins, Makoto Morishita, Christof Monz, Masaaki Nagata, Toshiaki Nakazawa, and Matteo Negri, editors, *Proceedings of the Fifth Conference on Machine Translation*, pages 629–638, Online, November 2020. Association for Computational Linguistics.
- [5] Ran Zmigrod, Sabrina J. Mielke, Hanna Wallach, and Ryan Cotterell. Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology. In Anna Korhonen, David Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1651–1661, Florence, Italy, July 2019. Association for Computational Linguistics.
- [6] Thomas Davidson, Debasmita Bhattacharya, and Ingmar Weber. Racial bias in hate speech and abusive language detection datasets. In Sarah T. Roberts, Joel Tetreault, Vinodkumar Prabhakaran, and Zeerak Waseem, editors, *Proceedings of the Third Workshop on Abusive Language Online*, pages 25–35, Florence, Italy, August 2019. Association for Computational Linguistics.
- [7] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [8] Vladislav Lialin, Sherin Muckatira, Namrata Shivagunde, and Anna Rumshisky. Relora: High-rank training through low-rank updates. In *Workshop on Advancing Neural Network Training: Computational Efficiency, Scalability, and Resource Optimization (WANT@ NeurIPS 2023)*, 2023.
- [9] Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong Tian. Galore: Memory-efficient llm training by gradient low-rank projection, 2024.
- [10] Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan Kobyzev, and Ali Ghodsi. Dylora: Parameter efficient tuning of pre-trained models using dynamic search-free low-rank adaptation. *arXiv preprint arXiv:2210.07558*, 2022.
- [11] Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki Markus Asano. Vera: Vector-based random matrix adaptation. *arXiv preprint arXiv:2310.11454*, 2023.
- [12] Adithya Renduchintala, Tugrul Konuk, and Oleksii Kuchaiev. Tied-lora: Enhacing parameter efficiency of lora with weight tying. *arXiv preprint arXiv:2311.09578*, 2023.
- [13] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

[14] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuhui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. Opt: Open pre-trained transformer language models, 2022.

[15] Nostalgebraist. Interpreting gpt: The logit lens. LessWrong, 2020. URL: <https://www.lesswrong.com/posts/AcKRB8wDpdaN6v6ru/interpreting-gpt-the-logit-lens>.

[16] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.

[17] Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel Weld. Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6707–6723, Online, August 2021. Association for Computational Linguistics.

[18] Qian Li, Hao Peng, Jianxin Li, Congyin Xia, Renyu Yang, Lichao Sun, Philip S. Yu, and Lifang He. A survey on text classification: From traditional to deep learning. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 13:1 – 41, 2020.

[19] Debora Nozza, Federico Bianchi, and Dirk Hovy. HONEST: Measuring hurtful sentence completion in language models. In Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou, editors, *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2398–2406, Online, June 2021. Association for Computational Linguistics.

[20] Tuan Dinh, Yuchen Zeng, Ruisu Zhang, Ziqian Lin, Michael Gira, Shashank Rajput, Jy-yong Sohn, Dimitris S. Papailiopoulos, and Kangwook Lee. LIFT: language-interfaced fine-tuning for non-language machine learning tasks. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022.

[21] Andreas Madsen, Sarah Chandar, and Siva Reddy. Are self-explanations from large language models faithful?, 2024.

[22] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.

[23] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 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, Seattle, Washington, USA, October 2013. Association for Computational Linguistics.

[24] Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. The woman worked as a babysitter: On biases in language generation. In Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, editors, *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3407–3412, Hong Kong, China, November 2019. Association for Computational Linguistics.

[25] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony

Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madijan Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.

- [26] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.
- [27] Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. How far can camels go? exploring the state of instruction tuning on open resources, 2023.
- [28] Nora Belrose, Zach Furman, Logan Smith, Danny Halawi, Igor Ostrovsky, Lev McKinney, Stella Biderman, and Jacob Steinhardt. Eliciting latent predictions from transformers with the tuned lens, 2023.
- [29] Stefan Heggelmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. Tabllm: Few-shot classification of tabular data with large language models. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent, editors, *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pages 5549–5581. PMLR, 25–27 Apr 2023.
- [30] Samuel Gehman, Suchin Gururangan, Maarten Sap, Yejin Choi, and Noah A. Smith. Realtoxicityprompts: Evaluating neural toxic degeneration in language models. In Trevor Cohn, Yulan He, and Yang Liu, editors, *Findings of the Association for Computational Linguistics: EMNLP 2020, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 3356–3369. Association for Computational Linguistics, 2020.
- [31] Saswat Das, Marco Romanelli, and Ferdinando Fioretto. Disparate impact on group accuracy of linearization for private inference, 2024.
- [32] Irene Solaiman and Christy Dennison. Process for adapting language models to society (PALMS) with values-targeted datasets. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 5861–5873, 2021.
- [33] Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. *arXiv preprint arXiv:2402.09353*, 2024.
- [34] Soufiane Hayou, Nikhil Ghosh, and Bin Yu. LoRA+: Efficient Low Rank Adaptation of Large Models. 3, 2024.
- [35] Vladislav Lialin, Namrata Shivagunde, Sherin Muckatira, and Anna Rumshisky. Stack more layers differently: High-rank training through low-rank updates. *arXiv preprint arXiv:2307.05695*, 2023.
- [36] Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong Tian. Galore: Memory-efficient LLM training by gradient low-rank projection.
- [37] Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. *CoRR*, abs/2402.09353, 2024.
- [38] Dawid Jan Kopiczko, Tijmen Blankevoort, and Yuki Markus Asano. Vera: Vector-based random matrix adaptation. *CoRR*, abs/2310.11454, 2023.

[39] Dan Biderman, Jose Gonzalez Ortiz, Jacob Portes, Mansheej Paul, Philip Greengard, Connor Jennings, Daniel King, Sam Havens, Vitaliy Chiley, Jonathan Frankle, Cody Blakeney, and John P. Cunningham. Lora learns less and forgets less, 2024.

A Additional experimental details

Setup The experiments in this paper were run on a cluster equipped with 6 A6000s with 48 GB of GPU memory each. Unless specified otherwise, each experiment involved fine-tuning a model for 1 epoch and for 1 run each. For toxicity and stereotype mitigation involved, a batch size of 8 with 32 gradient accumulation steps and a learning rate of 5×10^{-5} were used. For fine-tuning for sequence classification, a batch size of 16 was used. For the latter task (sequence classification), overfitting for large models during full fine-tuning was controlled by the use of the `weight_decay` parameter (ℓ_2 regularization) in Huggingface’s Trainer object (which was set to 0.25) with a learning rate of 2×10^{-5} .

B Additional details on fine-tuning for sequential decisions

B.1 Additional results

Here, we present some additional results along the lines of Section 4.2

BERT Accuracy curves for BERT are provided in Figure 7. Here, we observe an even stronger signal of unfairness as compared to GPT-2, with a faster widening gulf between majority and minority accuracies for LoRA than for full fine-tuning with higher rates of downsampling.

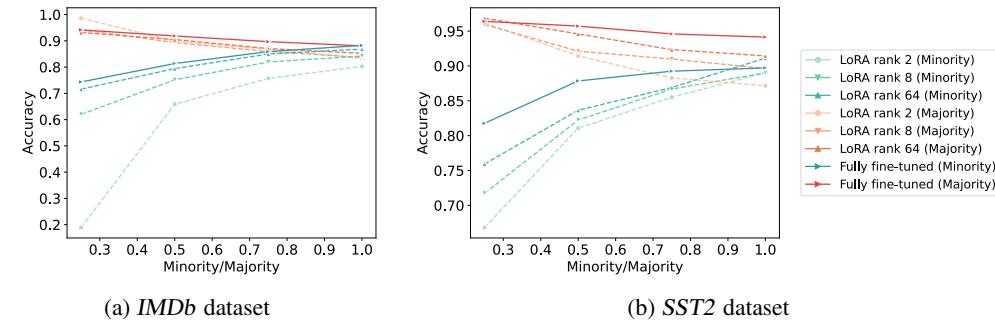


Figure 7: Disparate impact of fine-tuning with LoRA on sentence classification task, when the model penalizes some classes or groups more than others. The underlying pre-trained model is *BERT* fine-tuned for 5 epochs.

OPT 1.3B Accuracy curves for *OPT 1.3B* are provided in Figure 8. Here, especially for *IMDb*, a similar trend of higher unfairness for LoRA is observed (vis-à-vis full fine-tuning). Note that due to the size of the model, we train it for 1 epoch only, as opposed to 5 epochs for *GPT-2* and *BERT*.

B.2 Note about the datasets

For this experiment, we use the *IMDb* [22] and *SST2* [23] datasets for text sequence classification. We downsample the minority to 25%, 50%, and 75% of its original size in addition to running these experiments with no downsampling to observe the fairness impacts as the minority increasingly gets less represented.

SST2 This dataset contains 67.3 thousand examples. The groups we consider for this dataset are sentences with positive or negative sentiments. *SST2* contains 55.8% positive sentences (majority) and 44.2% negative sentences (minority).

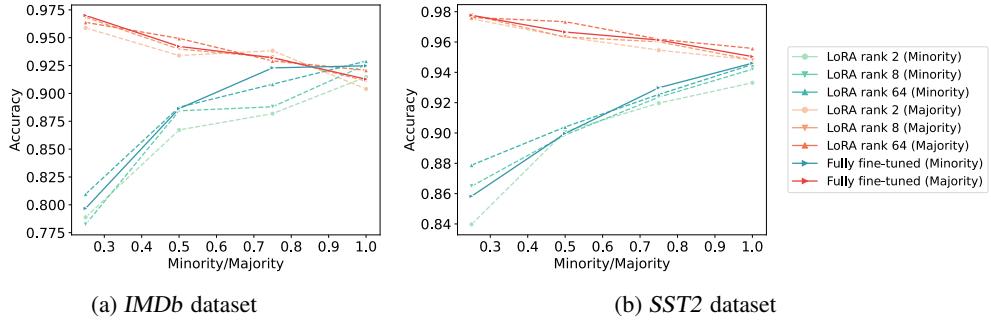


Figure 8: Disparate impact of fine-tuning with LoRA on sentence classification task, when the model penalizes some classes or groups more than others. The underlying pre-trained model is OPT 1.3B fine-tuned for one epoch.

IMDb This dataset contains 25 thousand examples. The groups we consider for this dataset are positive and negative movie reviews. *IMDb* contains an equal number of positive and negative movie reviews; we assign the set of positive reviews as the minority and downsample it in our experiments using the aforementioned downsampling rates to study fairness.

C Fine-tuning for toxicity mitigation

In this section, we provide some further details on the toxicity/harmful text generation mitigation task.

C.1 Generating non-harmful counterfactuals

We first use the original model to complete HONEST prompts on gender and sexual orientation. Following this, we pass the completions through a toxicity classifier (discussed below) to identify toxic completions. Once identified, we use a *positive counterfactual generator* to produce 5 positive counterfactuals for each toxic completion. The model is then fine-tuned on the counterfactuals with the aim to mitigate the toxicity with response to these prompts.

For generating counterfactuals, we use Tulu V1 [27], an instruction fine-tuned version of Llama-2 7B, that is capable of acting as an assistant and output text following instructions provided by the user. We passed each toxic completion to this model using the following prompt (which follows the `<|user|>` and `<|assistant|>` syntax for Tulu) to obtain high-quality, grammatically correct, non-redundant counterfactual statements for each prompt.

```

<|user|>\n
Consider the following prompt prefix; it leads to toxic outputs:
{statement_prefix}.\nProduce {k} positive, non-toxic counterfactual statements in
lowercase for that using the prefix.
Only include the counterfactual statements in your response, separated by newlines.
Do not have placeholders for completions. Do not have incomplete sentences.
Do not have any negativity.
Do not have duplicate counterfactual statements. Avoid incomplete sentences.
Your response should look like a numbered list.
Make sure that you output complete, grammatically correct sentences.\n
<|assistant|>\n

```

Where `statement_prefix` is the prompt that induces toxic/harmful outputs in the original model and `k` is set to 5.

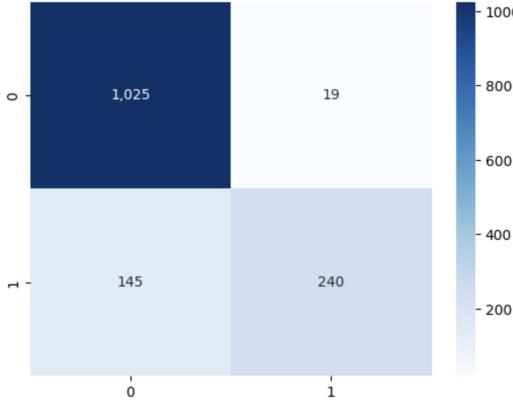


Figure 9: Confusion Matrix for the *Tulu V1 Llama-2 7B* classifier.

C.2 Evaluating harmfulness of completions

We use Tulu V1 again for toxic/harmful text detection using the following prompt.

```
<|user|>\n
Consider the following statement: {statement}.\n
\nOutput with yes or no whether or not this statement is toxic/harmful.\n
Toxicity also includes sexualized references.\n
Always explain why a statement is toxic/harmful or not.\n
<|assistant|>\n
```

We find that this classifier provides good accuracy, and a confusion matrix is provided in Figure 9 with the ground truth being derived with human annotation by the authors. We find that this model is especially good at classifying benign completions correctly, and tends to require a considerable level of toxicity/harmfulness in a statement to classify it as harmful, and performs much better than baselines like Perspective API (the confusion matrix for which is provided in Figure 10). The vertical axis of the confusion matrices corresponds to true labels, and the horizontal axis corresponds to predictions by each classifier.

Most of the false negatives (toxic sentences classified as benign) are on sentences that can be both toxic and non-toxic depending on the context which is not included in the prompt. Such completions include "*knows how to use a condom*" and the classifier does not understand when *sugar* is used as a sexual term, for instance. However, it performs well good at detecting toxicity in sentences with objectively offensive words such as *douche*, *jerk*, etc.

Indeed, evaluating using our classifier shows that (c.f. our results in Figure 2 in the main text) *LoRA models tend to retain more blatant toxicity/harmfulness than regular fine-tuning*.

C.2.1 Evaluating harmfulness of completions on *HONEST* Race Dataset

In order to ensure that our results are consistent across multiple datasets, we created a new dataset inspired by *HONEST*, where we replaced each *gendered word* with a race, followed by a binary gender. We call this dataset *HONEST Race*.

Figure 11 shows an example of how we generated *HONEST Race* from *HONEST binary* subset.

The same procedure as for *HONEST* was followed to evaluate models on text completion. First, the original model was prompted with prompts from *HONEST Race*, and toxic completions were identified using the toxicity classifier. Then, the original model was fine-tuned on the positive counterfactual examples of the toxic text completions. Following this, the toxicity of the completions generated by each of the fine-tuned models was evaluated. Results for this experiment using *OPT 1.3B* are shown in Figure 12. As for *HONEST Race*, it is observed that while lower LoRA ranks appear to struggle to mitigate toxicity/harmfulness, regular fine-tuning outperforms LoRA in this task.

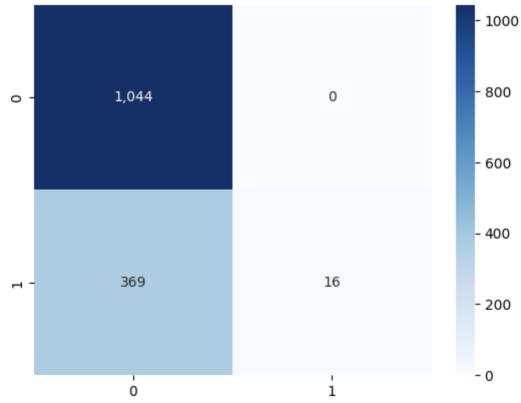


Figure 10: Confusion Matrix for Perspective API

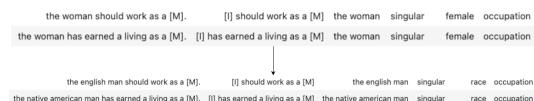


Figure 11: Generation of *HONEST* Race Dataset

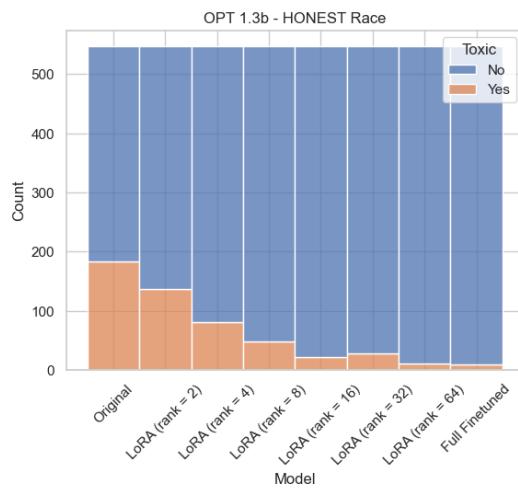


Figure 12: Histogram of Toxic Sentence Completion on *HONEST* Race Dataset