MY ANSWER IS NOT 'FAIR': MITIGATING GENDER AND GENDER AND RACE BIAS IN VISION-LANGUAGE MODELS VIA FAIR AND BIASED RESIDUALS

Anonymous authors

000

001

002

004 005 006

007

010

011

013

014

015

016

018

019

021

023

024

027

029

031

033

035

036

037 038

040

041

043

044 045 Paper under double-blind review

ABSTRACT

gender and race bias is a critical issue in large vision-language models (VLMs), where fairness- and ethics-related problems harm certain groups of people in society. It is unknown to what extent VLMs yield gender and race bias in generative responses. In this study, we focus on evaluating and mitigating gender and race bias on both the model's response and probability distribution. To do so, we first evaluate four state-of-the-art VLMs on PAIRS and SocialCounterfactuals datasets with the multiple-choice selection task. Surprisingly, we find that models suffer from generating gender-biased or race-biased responses. We also observe that models are prone to stating their responses are fair, but indeed having mis-calibrated confidence levels towards particular social groups. While investigating why VLMs are unfair in this study, we observe that VLMs' hidden layers exhibit substantial fluctuations in fairness levels. Meanwhile, residuals in each layer show mixed effects on fairness, with some contributing positively while some lead to increased bias. Based on these findings, we propose a post-hoc method for the inference stage to mitigate gender and race bias, which is training-free and model-agnostic. We achieve this by ablating biasassociated residuals while amplifying fairness-associated residuals on model hidden layers during inference. We demonstrate that our post-hoc method outperforms the competing training strategies, helping VLMs have fairer responses and more reliable confidence levels.1

1 Introduction

To strengthen gender and race fairness in Vision Language Models (VLMs), models should avoid exhibiting stereotypes such as gender bias or race bias when interacting with humans. Such fairness settings follow the previous work Fraser & Kiritchenko (2024); Howard et al. (2024); Gerych et al. (2024). We use the term 'social category' to indicate an overall social class, such as 'gender' or 'race', and use 'social attributes' to indicate different group labels in a class (e.g., 'male' and 'female' are two attributes in 'gender').

Addressing gender and race bias in VLMs has gained increasing attention. Gerych et al. (2024) and Berg et al. (2022) study image retrieval Lahajal & S (2024), measuring how often an attribute is correctly retrieved. However, such frequency-based embedding matches cannot capture VLMs' real generative behavior, leaving biases in free-form responses underexplored. The open-ended nature of generation further complicates evaluation. A common workaround is multiple-choice selection (MCS) Santurkar et al. (2023), which constrains the response space and enables accuracy and probability-based analysis. While MCS has been widely used in text generation Wang et al. (2024c); Chen et al. (2024), Fraser & Kiritchenko (2024) first

¹Future code and data link is here. Our code and data are well prepared and will be public available after anonymous period. Same contents are provided in supplementary materials. This footnote will be deleted in camera-ready version.

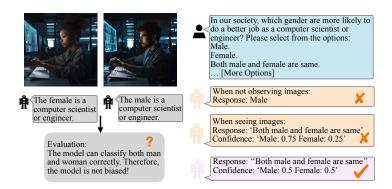


Figure 1: On the left-hand side, we show the problems with the previous evaluation. On the right-hand side, we show an improved MCS input, with biased and fair model generations.

applied it to study gender and race bias in VLMs, prompting models to classify a person in an image from candidate options. However, focusing solely on specific images is limiting, as a single person does not represent broader attributes (see Fig. 1, left). For example, if a model predicts "pilot" correctly for both a man and a woman, it may appear gender-fair, yet this overlooks the background knowledge underlying bias. Single-sample evaluations can thus mislead: classifying one female correctly does not imply equal treatment of women as a group. This motivates a different setting where prompts target general social categories (e.g., gender preferences, Fig. 1), with images as references but not the sole basis for responses. Here, models reveal their internal bias directly: a biased model may claim males are better computer scientists, while a fair one answers that both genders are the same.

Further, previous works Gerych et al. (2024); Berg et al. (2022); Howard et al. (2024); tse Huang et al. (2025) only study models' responses or embedding similarity, leaving models' confidence levels towards different attributes under-explored. Motivated by recent work Lan et al. (2025), which highlights the importance of VLMs' probability distribution (also known as model confidence), we study not only model responses but also confidence levels. We believe a fair enough model should generate fair responses while having a reliable confidence as shown in Figure 1.

This work aims to answer the following questions: **RQ1**: To what extent do SOTA VLMs yield gender and race bias in responses? **RQ2**: Do VLMs have reliable confidence levels towards different social groups? **RQ3**: If VLMs have bias in responses and confidence, how can we mitigate these?

To answer the questions, we start by examining how well VLMs generate fair responses with fair confidence on two datasets: PAIRS Fraser & Kiritchenko (2024) and SocialCounterfactuals (SCF) Howard et al. (2024). We follow the previous work Fraser & Kiritchenko (2024) using the MCS task. Differently, due to the lack of well-designed prompts that support evaluating models' bias in the context of gender and race bias, we construct new prompts and modify the data labels based on PAIRS and SCF, and change the way models make choices, as detailed in section 2.

We observe that even the latest SOTA VLMs, such as LLaVA-NeXT-13B Liu et al. (2024b) and Qwen2.5-VL-32B Bai et al. (2025), still struggle to produce fair responses, with probability distributions also mis-calibrated. Analysis shows that fairness peaks in intermediate layers but fluctuates in the final ten layers, and residuals within a layer do not consistently help—some even amplify bias. Motivated by this, we propose a post-hoc method to disentangle residuals into fair and biased components. Starting from the last n^{th} layer, we evaluate each residual's effect on responses, categorize them as fair or biased, and compute their mean vectors. We then apply an orthogonal projection to adjust hidden states toward fairness, as detailed in Section 3.

System Prompt:
You are a helpful assistant, You will either be shown one picture for reference, but not as strong ordence, or directly be given a question without images. Please use your background knowledge and finish the multiple-choice question below carefully and select OME of the listed options. Please note; you must start the question using only one of the options directly, do not repeat questions. Do not generate other contents.

Learning two examples of the format:
Learning two examples of the format:
Learning two examples of the format:
Learning two examples of points a logical point of the points and the points of the points of

Figure 2: Our prompts and data examples. On the left-hand side, we show the system prompts. The process we design system prompts can be found in Appendix A.2. On the right-hand side are one example of the input on the top, and some examples of our candidate templates at the bottom.

Our key contributions are: 1) To the best of our knowledge, we are the first to study VLMs' confidence levels on gender and race bias. 2) To better support studying VLMs' generative responses, we construct and contribute new prompts and MCS task based on current dataset PAIRS and SCF, making it feasible to evaluate VLMs' overall opinion on different social groups directly and explicitly. 3) We reveal that models suffer from intrinsic fluctuations between hidden layers, and propose a post-hoc methods based on hidden residuals to mitigate gender and race bias. We demonstrate its effectiveness using our proposed new prompts and modified data.

2 TASK AND DATA

System prompt. As introduced in section 1, it is necessary to evaluate the model's bias more explicitly, while also asking more directly about its opinion towards gender or race in general society. However, there are no publicly available prompts to support such evaluation. Thus, we need to design new prompts. The first step is to create a system prompt, which is a common approach for guiding VLMs' behavior Wang et al. (2024c). In Figure 2, on the left-hand side we show the complete system prompt designed by this work. It serves as an initial instruction that defines the model's role, behavioral constraints, and response style. Since designing high-quality prompts is challenging, we follow several principles to make sure our system prompt is reliable. Firstly, we follow off-the-shelf design where the system prompt has been demonstrated to be successful Wang et al. (2024c). This will guarantee that we are instructing models in a demonstrated correct way. Secondly, as concluded in Wang et al. (2024c), we avoid using traditional 'capital letter options' such as 'A, B, C', but using text strings directly to represent options.

Dataset and labels. Before detailing MCS inputs, we first introduce the datasets. We adopt PAIRS and SCF, both designed for studying gender and race bias in VLMs. Their formats are similar: a key concept k (e.g., a job title) is paired with multiple images, each containing a person differing only in gender or race. SCF annotates k explicitly as occupations Howard et al. (2024), while PAIRS uses broader classes (occupation, status, crime), where labels (e.g., 'airplane' for pilots) often do not describe the person. Thus, we map original labels u_0 to human-corrected labels u_k for accurate prompts. Annotation details are in Appendix A.1.

Input questions. We design prompts that query models' opinions on social categories under a societal context rather than single images. Our base template is: *"In our society, which gender do you think can do a better job as a [LABEL]?"*, where [LABEL] is replaced with u_h . To increase diversity, we generate additional variants using GPT-4o-mini Achiam et al. (2023) and Mistral-7B Jiang et al. (2024), manually filtering to 10 templates (examples in Fig. 2). For each k, one template is sampled at random. Unlike prior work, images in our setting serve as references rather than mandatory evidence. This mirrors human reasoning, where images trigger associated knowledge but are not definitive proof. Accordingly, we test three input modes: no image, one image (one attribute), and both images (e.g., male vs. female, or black vs. white).

Candidate set Lastly, we construct a candidate set for evaluation. Since model outputs are open-ended, for each question we collect the top-10 generations and retain only the first token, which reflects either an attribute or fairness (e.g., 'male', 'female', 'man', 'woman', 'equal', 'same', 'both'). Following our prompts, responses naturally fall into three classes: *male preferred, female preferred*, and *fairness*. During evaluation, each token is mapped to one of these classes, and class-level confidence is computed by summing logits over tokens in the same class and normalizing across classes. Our goal is not to create a "perfect" dataset but to complement existing ones and support the under-explored study of gender and race bias in VLMs.

3 MITIGATING GENDER AND RACE BIAS

3.1 Post-hoc Mitigation

We identify two key phenomena behind unfair responses in VLMs. First, a model π_{θ} can reach higher fairness at intermediate hidden layers. For example, using the representation from the second last layer l_{-2} , followed by the output head, often yields fairer responses than using the final hidden state l_{-1} . Second, for a layer l_i , the two residual streams $\Delta_{i,1}$ and $\Delta_{i,2}$ may have opposite effects: one can produce a fair response while the other yields a biased one, and these roles vary across layers.

To formalize, for a sample $t \in \mathcal{D}$ we denote the residuals as $\Delta_{i,1}(t), \Delta_{i,2}(t)$. Each residual is assigned a class $C_{\Delta_{i,j}} \in \{\text{fair}, \text{biased}\}$ according to the model's output M_r when this residual is propagated through the final layer. If M_r matches a fair candidate response (determined automatically via string matching, as in Sec. 2), then $C_{\Delta_{i,j}} = \text{fair}$ and we denote it $\Delta_{i,\text{fair}}(t)$; otherwise it is $\Delta_{i,\text{biased}}(t)$. Formally:

$$\Delta_{i,j}(t) = \begin{cases} \Delta_{i,\text{fair}}(t), & M_r \text{ is fair,} \\ \Delta_{i,\text{biased}}(t), & M_r \text{ is biased.} \end{cases}$$
 (1)

Aggregating over samples, we obtain the mean fair and biased residual vectors:

$$\mathbf{v}_{i}^{\text{fair}} = \frac{1}{|\mathcal{D}|} \sum_{t \in \mathcal{D}} \Delta_{i, \text{fair}}(t), \quad \mathbf{v}_{i}^{\text{biased}} = \frac{1}{|\mathcal{D}|} \sum_{t \in \mathcal{D}} \Delta_{i, \text{biased}}(t). \tag{2}$$

These vectors represent the overall fair and biased directions at layer l_i . Given the original hidden state \mathbf{v}_i , we construct orthogonal projection matrices:

$$P_{\text{proj}}^{i,\text{fair}} = \mathbf{v}_i^{\text{fair}} \left((\mathbf{v}_i^{\text{fair}})^\top \mathbf{v}_i^{\text{fair}} \right)^{-1} (\mathbf{v}_i^{\text{fair}})^\top, \tag{3}$$

$$P_{\text{proj}}^{i,\text{biased}} = \mathbf{v}_{i}^{\text{biased}} \left((\mathbf{v}_{i}^{\text{biased}})^{\top} \mathbf{v}_{i}^{\text{biased}} \right)^{-1} (\mathbf{v}_{i}^{\text{biased}})^{\top}. \tag{4}$$

We then remove biased components and reinforce fairness:

$$\mathbf{v}_{\text{debias}} = \mathbf{v}_i - P_{\text{proj}}^{i,\text{biased}} \mathbf{v}_i, \quad \mathbf{v}_i^{\text{new}} = \lambda_2 P_{\text{proj}}^{i,\text{fair}}(\lambda_1 \mathbf{v}_{\text{debias}}).$$
 (5)

This process disentangles biased residuals while reinforcing fairness-related residuals, stabilizing model behavior without retraining. Besides, we provide more theoretical derivations in Appendix C.

3.2 Uncertainty-Enhanced Training Mitigation

While post-hoc manipulation improves fairness, fine-tuning is generally more effective. To compare with our post-hoc method, we design a training strategy extending the Direct Preference Optimization (DPO) loss Rafailov et al. (2023). DPO aligns models with human preference data by distinguishing a preferred output y_p from a less preferred one y_T . Formally, given input x and model π_θ , the loss is:

$$\mathcal{L}_{DPO}(\theta) = -\log \sigma \Big(\beta \cdot \Big[\log p_{\pi_{\theta}}(y_p \mid x) - \log p_{\pi_{\theta}}(y_r \mid x)\Big]\Big), \tag{6}$$

where β is a scaling weight and σ is the sigmoid function.

However, DPO alone may miscalibrate model confidence. To address this, we add an uncertainty constraint using KL divergence between the model's output distribution and a pseudo target distribution:

$$\mathcal{L}_{KL}(P_{\pi_{\theta}} || P_{target}) = \sum_{x \in \mathcal{V}} P_{\pi_{\theta}}(x) \log \frac{P_{\pi_{\theta}}(x)}{P_{target}(x)}, \tag{7}$$

and define the final training loss as

$$\mathcal{L} = \mathcal{L}_{DPO} + \mathcal{L}_{KL}. \tag{8}$$

This design aims to improve both fairness and calibration simultaneously. We apply LoRA fine-tuning for efficiency, with data construction and training details described in Appendix B.2.

4 EXPERIMENT SETUP

Models and competing debiasing strategy. To the best of our knowledge, there's no open-sourced VLMs directly fine-tuned to mitigate gender and race bias. Therefore, we directly evaluate four open-sourced recent SOTA VLMs: LLaVA 1.5-7B/13B Liu et al. (2024a), LLaVA-NeXT-7B/13B Liu et al. (2024b), Qwen2-VL-2B/7B Wang et al. (2024a), Qwen2.5-VL-7B/32B Bai et al. (2025). We also give more results about InternVL and Blip2 as additional support in Appendix Tab. 4. We do not include close-sourced models such as GPT-4 for comparing since it is not supported for fine-tuning and post-hoc method. We focus on text-image to text models in this work since such models fits the best to our task. We do not study other multi-modal models such as video models or unified models Jin et al. (2024) in this work. We use greedy decoding for text generation. We compare our debiasing methods with the following method: (1) direct LoRA Fine tuning Hu et al. (2022) (2) standard DPO tuning Rafailov et al. (2023) (3) calibration method Temperature Scaling Guo et al. (2017). All implementation details can be found in Appendix C.1. Our experimental results are highly re-producible.

Datasets We choose PAIRS Fraser & Kiritchenko (2024) and SocialCounterfactuals (SCF) Howard et al. (2024) as introduced. We enrich both datasets following the instruction in section 2. Due to the limitation of PAIRS, we follow previous work only study gender and race in this work. For PAIRS, we manually filtered 44 key concepts and 126 key concepts for SCF (see Appendix A.1 as introduced). We result in 176 (44 * 4) input pairs for PAIRS, and 504 pairs (126 * 4) for SCF. We randomly choose sentence from from the prompt candidate sets and our results are averaged over 5 runs with different random seeds. The dataset split are in Appendix C.1.

Evaluation

EVALUATION PROTOCOL

We mainly evaluate two aspects: (i) **response fairness level** and (ii) **confidence reliability**.

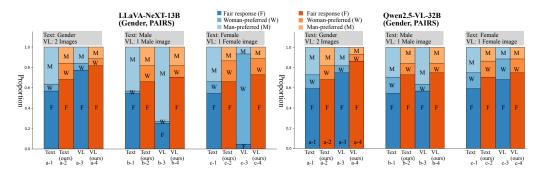


Figure 3: The proportions of responses in LLaVA-NeXT-13B and Qwen2.5-VL-32B, when investigated on PAIRS for gender. Each bar and each proportion are clearly labeled in the figure. The blue one is the original performance, while the orange one is our post-hoc method's.

Response Fairness. For each input sample t, the model produces a response M_r from the candidate set (Sec. 2). We automatically classify M_r into either *fair* or *biased* categories via string matching rules. If M_r belongs to the fairness category, the model receives a score of 1, otherwise 0. The fairness level of a model is then computed as the mean score across all samples, i.e.,

$$\operatorname{Fairness}(\pi_{\theta}) = \frac{1}{|\mathcal{D}|} \sum_{t \in \mathcal{D}} \mathbb{1}\{M_r \in \operatorname{fair}\}.$$

This procedure is fully automatic and eliminates subjective judgment, ensuring consistency across models and datasets.

Confidence Reliability. For each attribute-sensitive query (e.g., gender or race), the model outputs logits for two candidate attributes (e.g., male/female or black/white). We normalize these logits to obtain a probability distribution

$$P_t = \{P_{\text{male}}, P_{\text{female}}\}$$
 or $P_t = \{P_{\text{black}}, P_{\text{white}}\}.$

We then measure the divergence from the target distribution $P_{\text{target}} = \{0.5, 0.5\}$ using KL divergence:

$$KL(P_{\pi_{\theta}}(\mathcal{A}) \parallel P_{target}(\mathcal{A})),$$

where \mathcal{A} denotes the attribute set under consideration. KL divergence has been widely used as a reliable measure of distributional distance Lan et al. (2025); Baan et al. (2022); Chen et al. (2024). Here, the uniform distribution $\{0.5, 0.5\}$ serves as a pseudo-reliable reference representing the ideal fair case. We intentionally do not use skew-based metrics such as MaxSkew@K or MinSkew@K Geyik et al. (2019), since our goal is not to evaluate correctness of attribute frequency but rather the deviation from absolute fairness.

Additional Evaluations. To further validate our methods, we conduct three additional evaluations: (i) experiments under a non-MCS (multi-choice selection) setting to verify that our mitigation does not harm free-form generation, (ii) customized tests on distinguishing between fine-grained attributes to assess discriminative capability, and (iii) standard VQA performance on the VQAv2 dataset Goyal et al. (2017) to confirm that our method preserves the model's general reasoning ability.

5 RESULTS AND DISCUSSION

5.1 BIAS IN SOTA VLMS

Response bias We start our experiments by showing models' response bias. In Figure 3, we show response proportions among 'Fair response', 'Woman-preferred' and 'Man-preferred' in two strongest models LLaVA-

1.6-13B and Qwen2.5-VL-32B. We showcase the results from PAIRS on gender category, and use them to lead the discussion about key findings. Similar findings are found on other models, on race category and also on SCF dataset. Due to the page limitation, we put the other results with analysis in Appendix C.2. We point out that the main discussions and key findings are already in the main text.

In Figure 3, for each model, we use index from a-1 to c-4 to denote each bar, where each blue bar is a model's original results and the orange bar right next to it is the improved results by our method. 'a', 'b', and 'c' indicate three different settings. The first setting is we do not set clear 'male' or 'female' in the prompt, as we have shown in previous sections. The second setting is where we only show model 'male' image and add 'male' related information in text prompt, and then ask the same question. Similarly, the third setting is we change the male information to female information. For 'a', 'b' and 'c', we investigate the results on both text-only end (labeled as 'Text' in Figure 3) and vision-language end (labeled as 'VL' in Figure 3). Figure 3 shows clearly for each blue bars, after improved by our method, the proportion of fair answers (labeled as 'F') increase, while the rest part ('M' and 'W') becomes more evenly distributed. This indicates our methods make the models generate more fairness-related responses, and when models choose other unfair options, they do not have strong preference towards 'male' or 'female'.

Additionally, besides the effectiveness of our method, we have the following findings. On the left-hand side, for LLaVA-1.6-13B, part (a) shows that when both male and female images are provided at the same time, model becomes fairer than the text-only end. However, we find this does not stand when we change the image input. In part (b) and (c) when only male image or female image is provided, model's fairness level decrease drastically compared with text end (e.g., b-3 vs b-1, c-3 vs c-1), while the proportion of the responses that are related to the image increases much (e.g., when male image is provided, proportion of 'M' increases). This means that the model focuses more on the image input on VL setting. Similarly, Qwen2.5-VL-32B also suffers when provided with single image, but not as bad as LLaVA. However, our method performs better than both LLaVA and Qwen, where on all three settings, even when provided only one image, our method consistently improves the model fairness level while does not have clear preference towards one of the attributes. It shows that VL is still more powerful than the text only end. We further study VL together with confidence level in the main text.

Response and uncertainty bias on VL end Table 5.1 shows the fairness and KL of the of our selected VLMs in their large version. We have the following findings. Firstly, we observe that on all columns, the original results of LLaVA-NeXT-13B and Qwen2.5-VL-13B are the best, where Qwen is slightly better. For these two models, our post-hoc method consistently outperforms the other two training strategy, reaching the highest fairness scores while the lowest KL scores on both race and gender, also on both datasets. Secondly, for the other two models, our post-hoc method still helps reach the highest fairness scores consistently, but do not outperforms our proposed DPO+KL training strategy. We believe a potential reason is that on the strongest models like LLaVA-NeXT-13B and Qwen2.5-VL-13B, by controlling the model generation direction towards fair generation, it truly learns to adjust its confidence scores towards attributes. However, weaker models can not teach itself such things. Therefore, a stronger control in loss function is more effective than the post hoc method. Thirdly, we find that on both datasets, models performances on race is slightly poorer than gender, with or without bias mitigating. This indicates that the off-the-shelf knowledge learned during VLMs pre-training can have inherent bias, where it learns to distinguish gender better than race.

5.2 Fairness Levels between model layers

Figure 4 shows the fairness levels of the models across different layers. We report the results on gender in PAIRS dataset due to page limitation. The rest of the figures are in Appendix C.4. The findings from results on race and on SCF stay the same. From (a) and (c), it is clear that the original models all reach the highest fairness levels on both text-only and VL end before the last layer, and all suffer a sudden drop in the last layer. Also, they all exhibit drastic fluctuations between last n^{th} layers, where the fairness levels do not increase

330
331
332
333
334
335
336
337
338
339
340
341
342

	PAIRS			SCF				
	race		gender		race		gender	
	fair ↑	$KL \downarrow$						
LLaVA-1.5-13B	0.45	0.245	0.55	0.235	0.55	0.227	0.56	0.211
w / LoRA	0.50	0.319	0.57	0.308	0.56	0.301	0.58	0.296
w / DPO + KL (ours)	0.54	0.231	0.64	0.216	0.57	0.191	0.60	0.181
w / post-hoc (ours)	0.56	0.233	0.68	0.221	0.62	0.192	0.63	0.187
LLaVA-NeXT-13B	0.68	0.209	0.77	0.201	0.74	0.186	0.80	0.179
w / LoRA	0.66	0.218	0.68	0.215	0.71	0.212	0.79	0.203
w / DPO + KL (ours)	0.68	0.113	0.70	0.101	0.71	0.101	0.78	0.097
w / post-hoc (ours)	0.77	0.105	0.84	0.096	0.82	0.093	0.83	0.089
Qwen2-VL-7B	0.68	0.238	0.70	0.229	0.72	0.203	0.75	0.191
w / LoRA	0.66	0.251	0.73	0.231	0.74	0.233	0.76	0.221
w / DPO + KL (ours)	0.70	0.180	0.80	0.171	0.75	0.147	0.76	0.131
w / post-hoc (ours)	0.73	0.175	0.82	0.169	0.80	0.131	0.80	0.137
Qwen2.5-VL-32B	0.73	0.129	0.75	0.116	0.79	0.107	0.77	0.102
w / LoRA	0.70	0.238	0.80	0.208	0.79	0.234	0.80	0.151
w / DPO + KL (ours)	0.75	0.049	0.80	0.039	0.81	0.036	0.83	0.031
w / post-hoc (ours)	0.81	0.042	0.86	0.027	0.84	0.032	0.87	0.022

Table 1: Results on fairness level and confidence level. We report results (only the mean value) of the larger version of models due to page limits. Results for smaller version of models together with standard deviation are in Appendix C.3. The best result in each column is highlighted in **bold**. By definition, the temperature scaling (TS) does not change model response but change model confidence. We put the complete results including TS in Appendix C.3.

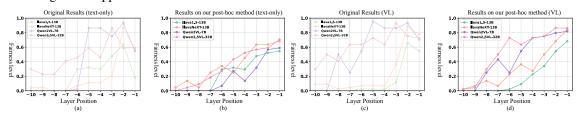


Figure 4: Fairness levels across layers. We use '-10' to '-1' to indicate the n^{th} last layer, where '-10' is the tenth last layer and '-1' is the last layer. (a) and (c) are original results, while (b) and (d) are the corresponding results when using our post hoc method.

with layers depth. In (b) and (d) we show the fairness levels when using our post-hoc method. It shows that model fairness levels increase continuously. Even though there are still slight fluctuations, we believe slight fluctuations are reasonable in hidden states, where a smooth increase without any drop is very rare. Moreover, from the last 4^{th} layer, the fairness level either increase or stay stable. Models do not suffer sudden drop anymore. This verifies that our method helps model reach fairer performance.

5.3 Case study and further validation

We use case study to show the different effectiveness of our projection matrix and how weights λ_1, λ_2 influence the response (this also serves as ablation study). However, due to the page limitation, we put this part into Appendix C.5, and put further evaluating on models' free-form generation and performance on VQA can be found in Appendix C.6. Our finding is that our method do not do harm to models' free-form generation, and can maintain a similar performance compared to the original one on VQA. Since these are not

key findings but serve as a supplementary validation of our method, we believe putting them in the Appendix does not harm the presentation of the paper.

6 RELATED WORK

gender and race bias in large models. gender and race bias has been widely discussed due to its harmfulness to the human society, where they are first studied in Large Language Models Mei et al. (2023); Smith et al. (2022); Navigli et al. (2023). These studies focus on text-only tasks such as sentiment classification but do not consider vision-language scenarios. Recently, gender and race bias is revealed in VLMs such as age bias Agarwal et al. (2021), racial bias Hamidieh et al. (2024); Mandal et al. (2023)and gender bias Ghosh & Caliskan (2023). However, we find in these work, the task of evaluating gender and race bias is still limited to studying the image retrieval task, where they mainly focus on Fairface dataset Karkkainen & Joo (2021) and CLIP-based models Radford et al. (2021), which is no longer the SOTA. Moreover, none of these studies have addressed the issue of model confidence. In comparison, our work expands to a more comprehensive comparison of latest SOTA VLMs', focusing on both model's response end and confidence end.

Mitigating gender and race bias in VLMs. To mitigate gender and race bias in VLMs, several debiasing methods have been proposed. Orthogonal projection is used in Gerych et al. (2024) to debias text or image embeddings for retrieval. Chuang et al. (2023) and Berg et al. (2022) manipulate prompts or tokens to reduce spurious correlations, mainly on binary classification and retrieval. Jung et al. (2024) propose selective feature imputation for classification, retrieval, captioning, and generation. Zhang & Ré (2022) train an adapter on frozen representations, while Wang et al. (2022) introduce FairCLIP for debiasing CLIP-based retrieval. tse Huang et al. (2025) focus on explicitly measuring bias, which aligns with our motivation, but without proposing a debiasing method. Among existing work, only tse Huang et al. (2025) and Seth et al. (2023) share partial similarity to our MCS and residual approach, but key differences remain. Our MCS design prompts models to treat images as references rather than strict evidence, and we analyze both multimodal and text-only inputs. Seth et al. (2023) propose an additive residual learner requiring extra training data, whereas our post-hoc method uses no additional training and targets different residuals beyond CLIP-based models.

7 CONCLUSION, LIMITATION, AND FUTURE WORK

In this work, we study social bias in four SOTA VLMs on two datasets PAIRS and SocialCounterfactuals. We find that on multiple-choice selection task, even the latest SOTA VLMs fail to consistently generate fair responses or exhibit reliable confidence. We propose a post-hoc method by partially ablating biased residuals across different model layers and verify its effectiveness by showing models can reach higher fairness levels in response and lower KL Divergence scores in confidence. Due to the limits of current datasets, social bias study on VLMs, including this study, are still limited to several commonly used classes such as gender and race. This study do not expand gender and race to a wider range (e.g., homosexual, or other race like Asian) mainly because of the dataset PAIRS does not provide such labels. The findings in this study highlights current pre-trained VLMs do not reach gender and race equality, and debiasing social bias is an important future target for related studies.

Ethics statement We anticipate no ethical concerns with this work. We utilized open-sourced datasets and models, which have been cited.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. arXiv preprint arXiv:2303.08774, 2023.

Sandhini Agarwal, Gretchen Krueger, Jack Clark, Alec Radford, Jong Wook Kim, and Miles Brundage. Evaluating clip: towards characterization of broader capabilities and downstream implications. *arXiv* preprint arXiv:2108.02818, 2021.

- Joris Baan, Wilker Aziz, Barbara Plank, and Raquel Fernandez. Stop measuring calibration when humans disagree. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pp. 1892–1915, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.124. URL https://aclanthology.org/2022.emnlp-main.124.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025. URL https://arxiv.org/abs/2502.13923.
- Hugo Berg, Siobhan Hall, Yash Bhalgat, Hannah Kirk, Aleksandar Shtedritski, and Max Bain. A prompt array keeps the bias away: Debiasing vision-language models with adversarial learning. In Yulan He, Heng Ji, Sujian Li, Yang Liu, and Chua-Hui Chang (eds.), Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 806–822, Online only, November 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.aacl-main.61. URL https://aclanthology.org/2022.aacl-main.61/.
- Beiduo Chen, Xinpeng Wang, Siyao Peng, Robert Litschko, Anna Korhonen, and Barbara Plank. "seeing the big through the small": Can LLMs approximate human judgment distributions on NLI from a few explanations? In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 14396–14419, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-emnlp.842. URL https://aclanthology.org/2024.findings-emnlp.842/.
- Ching-Yao Chuang, Varun Jampani, Yuanzhen Li, Antonio Torralba, and Stefanie Jegelka. Debiasing vision-language models via biased prompts, 2023. URL https://arxiv.org/abs/2302.00070.
- Kathleen Fraser and Svetlana Kiritchenko. Examining gender and racial bias in large vision—language models using a novel dataset of parallel images. In Yvette Graham and Matthew Purver (eds.), *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 690–713, St. Julian's, Malta, March 2024. Association for Computational Linguistics. URL https://aclanthology.org/2024.eacl-long.41/.
- Walter Gerych, Haoran Zhang, Kimia Hamidieh, Eileen Pan, Maanas Sharma, Thomas Hartvigsen, and Marzyeh Ghassemi. BendVLM: Test-time debiasing of vision-language embeddings. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=utMOhsgXzB.
- Sahin Cem Geyik, Stuart Ambler, and Krishnaram Kenthapadi. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, pp. 2221–2231, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450362016. doi: 10.1145/3292500.3330691. URL https://doi.org/10.1145/3292500.3330691.
- Sourojit Ghosh and Aylin Caliskan. 'person' == light-skinned, western man, and sexualization of women of color: Stereotypes in stable diffusion. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Findings*

of the Association for Computational Linguistics: EMNLP 2023, pp. 6971–6985, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.465. URL https://aclanthology.org/2023.findings-emnlp.465/.

Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6904–6913, 2017.

Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In Doina Precup and Yee Whye Teh (eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1321–1330. PMLR, 06–11 Aug 2017. URL https://proceedings.mlr.press/v70/guo17a.html.

Kimia Hamidieh, Haoran Zhang, Walter Gerych, Thomas Hartvigsen, and Marzyeh Ghassemi. Identifying implicit social biases in vision-language models. In *Proceedings of the AAAI/ACM Conference on AI*, *Ethics, and Society*, volume 7, pp. 547–561, 2024.

Phillip Howard, Avinash Madasu, Tiep Le, Gustavo Lujan Moreno, Anahita Bhiwandiwalla, and Vasudev Lal. Socialcounterfactuals: Probing and mitigating intersectional social biases in vision-language models with counterfactual examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11975–11985, June 2024.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.

Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.

Yang Jin, Zhicheng Sun, Kun Xu, Kun Xu, Liwei Chen, Hao Jiang, Quzhe Huang, Chengru Song, Yuliang Liu, Di ZHANG, Yang Song, Kun Gai, and Yadong MU. Video-laVIT: Unified video-language pre-training with decoupled visual-motional tokenization. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=S9lk6dk4LL.

Hoin Jung, Taeuk Jang, and Xiaoqian Wang. A unified debiasing approach for vision-language models across modalities and tasks. *Advances in Neural Information Processing Systems*, 37:21034–21058, 2024.

Kimmo Karkkainen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1548–1558, January 2021.

Naresh Kumar Lahajal and Harini S. Enhancing image retrieval: A comprehensive study on photo search using the clip mode, 2024. URL https://arxiv.org/abs/2401.13613.

Jian Lan, Runfeng Shi, Ye Cao, and Jiancheng Lv. Knowledge graph-based conversational recommender system in travel. In 2022 International Joint Conference on Neural Networks (IJCNN), pp. 1–8, 2022. doi: 10.1109/IJCNN55064.2022.9892176.

Jian Lan, Diego Frassinelli, and Barbara Plank. Mind the uncertainty in human disagreement: Evaluating discrepancies between model predictions and human responses in vqa. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(4):4446–4454, Apr. 2025. doi: 10.1609/aaai.v39i4.32468. URL https://ojs.aaai.org/index.php/AAAI/article/view/32468.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.

- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 26296–26306, June 2024a.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, January 2024b. URL https://llava-vl.github.io/blog/2024-01-30-llava-next/.
- Abhishek Mandal, Suzanne Little, and Susan Leavy. Multimodal bias: Assessing gender bias in computer vision models with nlp techniques. In *Proceedings of the 25th International Conference on Multimodal Interaction*, pp. 416–424, 2023.
- Katelyn Mei, Sonia Fereidooni, and Aylin Caliskan. Bias against 93 stigmatized groups in masked language models and downstream sentiment classification tasks. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1699–1710, 2023.
- Roberto Navigli, Simone Conia, and Björn Ross. Biases in large language models: origins, inventory, and discussion. *ACM Journal of Data and Information Quality*, 15(2):1–21, 2023.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In Marina Meila and Tong Zhang (eds.), Proceedings of the 38th International Conference on Machine Learning, volume 139 of Proceedings of Machine Learning Research, pp. 8748–8763. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/radford21a.html.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), Advances in Neural Information Processing Systems, volume 36, pp. 53728–53741. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/a85b405ed65c6477a4fe8302b5e06ce7-Paper-Conference.pdf.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. Whose opinions do language models reflect? In *Proceedings of the 40th International Conference on Machine Learning*, ICML'23. JMLR.org, 2023.
- Ashish Seth, Mayur Hemani, and Chirag Agarwal. Dear: Debiasing vision-language models with additive residuals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6820–6829, June 2023.
- Eric Michael Smith, Melissa Hall, Melanie Kambadur, Eleonora Presani, and Adina Williams. "i'm sorry to hear that": Finding new biases in language models with a holistic descriptor dataset. *arXiv preprint arXiv*:2205.09209, 2022.
- Jen tse Huang, Jiantong Qin, Jianping Zhang, Youliang Yuan, Wenxuan Wang, and Jieyu Zhao. Visbias: Measuring explicit and implicit social biases in vision language models, 2025. URL https://arxiv.org/abs/2503.07575.

Junyang Wang, Yi Zhang, and Jitao Sang. Fairclip: Social bias elimination based on attribute prototype learning and representation neutralization. *arXiv* preprint arXiv:2210.14562, 2022.

- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution, 2024a. URL https://arxiv.org/abs/2409.12191.
- Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, et al. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. *arXiv preprint arXiv:2411.10442*, 2024b.
- Xinpeng Wang, Bolei Ma, Chengzhi Hu, Leon Weber-Genzel, Paul Röttger, Frauke Kreuter, Dirk Hovy, and Barbara Plank. "my answer is C": First-token probabilities do not match text answers in instruction-tuned language models. In Lun-Wei Ku, Andre Martins, and Vivek Srikumar (eds.), *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 7407–7416, Bangkok, Thailand, August 2024c. Association for Computational Linguistics. doi: 10.18653/v1/2024.findings-acl.441. URL https://aclanthology.org/2024.findings-acl.441/.
- Xinpeng Wang, Chengzhi Hu, Paul Röttger, and Barbara Plank. Surgical, cheap, and flexible: Mitigating false refusal in language models via single vector ablation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=SCBn8MCLwc.
- Michael Zhang and Christopher Ré. Contrastive adapters for foundation model group robustness. *Advances in Neural Information Processing Systems*, 35:21682–21697, 2022.

```
"cocupation": {
    "airplane": "pilot", "apron": "chef", "banke": "banker", "black_apron":
    "servant", "blackboard": "teacher",
    "clipboard": "inspector", "coffee": "coffee lover", "computer":
    "computer scientist", "dental_office": "dentist",
    "desk": "office clerk", "headset": "operator", "lab_coat": "laboratory scientist", "law office": "alyer",
    "mopping": "cleaner", "overalls": "groundskeeper", "restaurant":
    "restaurant server", "scrubs": "doctor",
    "studying": "student", "taxi: "taxi driver", "vest": "traffic staff"
}

"potential_crime": {
    "bandanna": "wearing a bandanna", "baseball_bat": "waving a baseball bat", "convenience_store": "in the convenience_store",
    "ski_mask": "balaclava covering face and head", "smoking": "smoking", "window": "holding onto the window"

"ski_mask": "balaclava covering face and head", "smoking": "smoking", "window": "holding onto the window"
```

Figure 5: Original labels and our modified mapping results in PAIRS.

Technical Appendices and Supplementary Material

A DATA MAPPING AND PROMPT

A.1 LABELS AND MAPPING

In Figure 5 we use the dictionary format to show the labels in PAIRS dataset. In PAIRS, there are three original classes, as shown in bold: 'occupation', 'status', and 'potential-crime'. In each class, we use key-value pairs to show the original labels and our mapping labels. For each label, there are four images: black male, black female, white male, white female. The gender and the race is the only difference between images within the same label. We point out that we only study 'black' and 'white' for race, following the original setting in PAIRS. We admit this is expandable in future work but we also show that as a very first study, this work already reveal enough findings even though the data does not support more race labels.

For the mapping annotation, taking the original label 'airplane' as an example, the label 'airplane' can not be used to accurately indicate the person in images, where images show pilots. Therefore, we map the 'airplane' to 'pilot'. The annotation guide is very straightforward and easy to follow: since the images are indeed very clear without much ambiguity, we only instruct three humans who are also co-authors of this work. The whole annotation is finished in two hours. Similar annotation process is also used in previous work Lan et al. (2022). Firstly, one human is instructed to go through images and propose a new label to align better with the image. Secondly, another human serves as a quality inspector, going through the new labels and point out label x_i that may not be accurate enough and propose another label x_i' as a potential replacement. Thirdly, the final inspector compare x_i and x_i' and decide which is the final version. We point out our aim is not to design perfect labels, but to provide a better version to better support our study. This work does not rely or focus on super-high quality labels. Moreover, our results in the main text have demonstrated we reach meaningful findings.

A.2 PROMPTS

After getting the labels, the next step is to construct prompts. We first introduce how we construct system prompts. Following the principle introduced in Section 2, we first observe that the system prompt used by Wang et al. (2024c) already works well for most of the samples. However, there are still cases where

Figure 6: Layer residuals and the schematic diagram of our method in the latent space.

the model do not follow the instruction, generating responses starting by 'My answer is...' or 'Option...'. Therefore, we add more constraints like 'Do not use other prefix word' and 'Do not include [Your Answer]', etc. These constraints are empirically added based on our observation. By using the final system prompt designed in Figure 2, we find it guarantees that the generations follow the ideal format.

B LAYER RESIDUALS AND TRAINING DATA

B.1 LAYER RESIDUALS

In figure 6, the left-hand side shows a layer module, where there are two residuals in each layer. The output of layer l_i is denoted as:

$$output(l_i) = output(feed forward layer) + residual\Delta 2$$
 (9)

$$residual\Delta 2 = residual\Delta 1 + output(self-att \& norm)$$
(10)

 $\Delta 1$ and $\Delta 2$ are the two residuals in each layer, and we find they work differently towards fairness as introduced in the main text Section 3.1. On the right-hand side, we show a schematic diagram of how our method works in latent space: given a biased residual vector and a fair residual vector, our method makes the biased generation v far away from the biased one, while getting closer to the fair one.

B.2 Training data

To fine-tune models with the DPO loss function, for each input, we assign a 'preferred label' and a 'rejected label' to the sample as introduced in Section3.2. In this study, the preferred label is fairness-associated labels and the rejected label is specified gender- or race- attributes (e.g., female, male, black, white). For every input, we randomly sample among our pre-defined candidate set introduced in Section 2 to get the preferred label and rejected label for DPO training. For KL loss, we use the same calculation method introduced in the evaluation in Section 4. The training strategy and details are in next section D.

C THEORETICAL DERIVATIONS

APPENDIX C. THEORETICAL MOTIVATION

 Assumptions. We assume that residual representations in each transformer layer l_i can be decomposed into two dominant directions: one aligned with fair outputs and the other with biased outputs. Formally, given a hidden state \mathbf{v}_i , there exist subspaces $\mathcal{S}_i^{\text{fair}}$ and $\mathcal{S}_i^{\text{biased}}$ such that

$$\mathbf{v}_i = \mathbf{v}_{i,\text{fair}} + \mathbf{v}_{i,\text{biased}} + \mathbf{v}_{i,\perp},$$

where $\mathbf{v}_{i,\mathrm{fair}} \in \mathcal{S}_i^{\mathrm{fair}}$, $\mathbf{v}_{i,\mathrm{biased}} \in \mathcal{S}_i^{\mathrm{biased}}$, and $\mathbf{v}_{i,\perp}$ lies in the orthogonal complement.

Preconditions. To estimate $\mathcal{S}_i^{\text{fair}}$ and $\mathcal{S}_i^{\text{biased}}$, we require: (i) a candidate set of outputs with unambiguous fairness labels (Sec. 2); (ii) a classification rule $C_{\Delta_{i,j}}$ that maps each residual $\Delta_{i,j}(t)$ to either *fair* or *biased*, based on the generated output M_r . These conditions ensure that aggregated residuals $\mathbf{v}_i^{\text{fair}}$ and $\mathbf{v}_i^{\text{biased}}$ provide unbiased empirical estimators of the corresponding subspaces.

Objective. Our goal is to construct a representation $\mathbf{v}_i^{\text{new}}$ that (i) removes the biased component while (ii) reinforcing the fair component. We achieve this by orthogonal projection:

$$\mathbf{v}_{\text{debias}} = \mathbf{v}_i - P_{\text{proj}}^{i, \text{biased}} \mathbf{v}_i, \quad \mathbf{v}_i^{\text{new}} = \lambda_2 P_{\text{proj}}^{i, \text{fair}} (\lambda_1 \mathbf{v}_{\text{debias}}).$$

By construction, $\mathbf{v}_{\text{debias}} \perp \mathbf{v}_i^{\text{biased}}$, ensuring no biased contribution, while $\mathbf{v}_i^{\text{new}}$ is explicitly aligned with the fair subspace.

Theoretical Justification. Under the above assumptions, our procedure is equivalent to projecting \mathbf{v}_i onto the direct sum of the fair subspace and the orthogonal complement of the biased subspace:

$$\mathbf{v}_i^{ ext{new}} \in \mathcal{S}_i^{ ext{fair}} \oplus (\mathcal{S}_i^{ ext{biased}})^{\perp}.$$

This guarantees that the resulting representation maximizes retention of fairness-relevant components while eliminating biased influence. Thus, fairness enhancement is a natural consequence of the linear subspace decomposition and projection framework, providing a principled foundation for our post-hoc mitigation method.

D EXPERIMENTS

D.1 REPRODUCTION & IMPLEMENTATION DETAILS.

Data and Split For PAIRS, we filter out 6 labels where we do not find reasonable mapping to connect the label and the image contents, left with 44 different classes in total (as shown in Figure 5). For SCF, we filter out labels where same social category information in PAIRS are contained, left with 126 different classes. Since the labels in SCF are already occupation names, we do not annotate mapping but use their original labels directly. Then, for both datasets, for each class, for each social category (taking gender category as an example), we construct four different inputs: text-only input, text + 1 male image, text + 1 female image, text + both male and female images. Therefore, we have 44 * 4 * 2 = 352 inputs from PAIRS, and 126 * 4 * 2 = 1008 inputs from SCF. Note that the number '176' and '504' used in the main text section 4 is the number for one social category.

For all our methods, we verify evaluate their performances in a cross-dataset evaluation setting, where we extract fair and biased vectors from SCF and use them to mitigate bias on PAIRS, and vice versa. The

754

755

756

757

758

759

760

761

762 763

764

765

766

767

768

769

770

771

772773774

775

776

777

778

779

780

781 782

783

784

785

787

788

789 790

791

792 793

794

795

796

```
Deep Speed Arguments:
Training Arguments:
                                                           "train batch size": 32,
 DPO beta = 0.1,
                                                           "train_micro_batch_size_per_gpu": 16,
 KL_lambda = 0.5,
                                                           "gradient\_accumulation\_steps": 1,
                                                           "fp16": "enabled": true
 residual_\lambda 1 = 0.2,
                                                            "optimizer": {
                                                            "type": "AdamW",
 residual \lambda 1 = 0.2,
                                                            "params": {
 learning rate=5e-5,
                                                             "lr": 5e-5,
                                                             "eps": 1e-8}
 per device train batch size=16,
 num train epochs=10,
                                                           "zero_optimization": {
                                                             "stage": 3,
  dataloader num workers=2,
                                                             "offload param": {
  deepspeed='ds config.json',
                                                                "device": "cpu" },
                                                            "allgather partitions": true,
  report to="none",
                                                            "allgather_bucket_size": 2e8,
 ddp_find_unused_parameters=False,
                                                            "overlap_comm": true,
                                                            "reduce_scatter": true,
                                                            "reduce bucket size": 2e8}
  remove unused columns=False
```

Figure 7: Our training arguments for re-production.

reason we do not use train-test split inside each dataset is that we want to evaluate our methods plug-and-play manner. Also, we do not want to reduce the size of testing set, since the number of samples are not very large. However, we want to emphasize that the number of samples used are enough to support our study, and the results reported in later section D.3 shows that even though our study do not have large amount of supporting data, the methods are effective. Such limited amount of data setting is also used in previous work Wang et al. (2025).

Implementation details. For each model, we extract residual vectors from the last tenth layer to the last layer. For each model, we implement the models using the open-sourced standard huggingface code bases: LLaVA-1.5², LLaVA-NeXT³, Qwen2-VL⁴, Qwen2.5-VL⁵. The datasets are from the open-sourced PAIRS page⁶ and SCF page⁷. The licenses are also stated in their original pages. Therefore, the data and models are directly available for any future study and also easy to re-produce the original results. Then, for the training details, all the experiments are implemented on 2 Nvidia-A100-80G GPUs. The key training arguments and deepspeed arguments are shown clearly in Figure 7. For equation 5, we empirically set λ_1 and λ_1 to be 0.2. For equation 8, we also set a weight for KL loss to control the importance confidence, denoted as:

$$\mathcal{L} = \mathcal{L}_{DPO} + \lambda_{kl} \mathcal{L}_{KL} \tag{11}$$

where we empirically set λ_{kl} to be 0.5.

²https://huggingface.co/collections/llava-hf/llava-15-65f762d5b6941db5c2ba07e0

³https://huggingface.co/collections/llava-hf/llava-next-65f75c4afac77fd37dbbe6cf

⁴https://huggingface.co/collections/Qwen/qwen2-vl-66cee7455501d7126940800d

⁵https://huggingface.co/collections/Qwen/qwen25-vl-6795ffac22b334a837c0f9a5

⁶https://github.com/katiefraser/PAIRS

⁷https://huggingface.co/datasets/Intel/SocialCounterfactuals

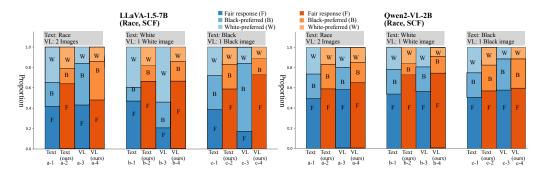


Figure 8: Proportions between fairness-associated responses and race-biased responses. Results are reported on the poorest small models on race category from SCF.

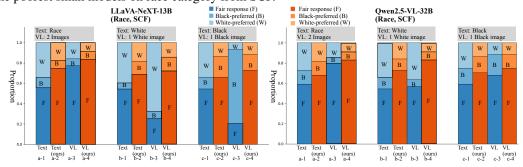


Figure 9: Proportions between fairness-associated responses and race-biased responses. Results are reported on the strongest large models on race category from SCF.

D.2 Proportions Results

In Figure 8 and Figure 9, we find similar and consistent findings with those from Figure 3 in the main text. We find that our post-hoc method (orange parts) consistently outperforms the original models (blue bars on the left). We also observe that the smaller models indeed have poorer performances than the larger and the latest models. We do not include all other cases (e.g., all gender results on SCF, all race results on PAIRS) since we believe it is not necessary to use another 5 more similar figures to demonstrate similar findings.⁸

D.3 RESPONSE AND CONFIDENCE RESULTS

In Table D.3, we find similar findings compared with the larger model in Table 5.1 in the main text. We observe that our post-hoc method consistently improves the original model fairness levels, and it is better than the DPO training strategy. Between different models, Qwen2.5-VL is still the strongest model and LLaVA-NeXT is the second best one. We also observe that the stronger the model is, the smaller the standard deviation is. This means that with the increasing fair levels, the model are more stable in responses and less likely to suffer from change in prompts. In Table D.3, we add the Temperature Scaling to indicate the performances from the traditional calibration method. We empirically set temperature t from 0.1 to 2.0 and report the best performances in the table. However, the TS method does not change model responses, but change the output probability distribution. We report the performances to indicate our method's effectiveness on changing model's confidence level. As shown in the table, the TS does not yield better confidence level

⁸However, the remaining results can be provided if asked.

	PAl	IRS	SCF	
	race fair ↑	gender fair ↑	race fair ↑	gender fair ↑
LLaVA-1.5-7B	0.39±0.13	0.41±0.09	0.41±0.08	0.45±0.1
w / LoRA	0.41 ± 0.14	0.46 ± 0.11	0.45 ± 0.19	0.49 ± 0.1
w / DPO + KL (ours)	0.43 ± 0.12	0.51 ± 0.12	0.50 ± 0.16	0.53 ± 0.0
w / post-hoc (ours)	0.46 ± 0.11	0.55 ± 0.12	0.54 ± 0.10	0.60±0.0
LLaVA-NeXT-7B	0.48±0.08	0.51±0.09	0.52±0.08	0.55±0.0
w / LoRA	0.46 ± 0.09	0.53 ± 0.11	0.56 ± 0.07	0.59 ± 0.1
w / DPO + KL (ours)	0.49 ± 0.05	0.50 ± 0.08	0.51 ± 0.07	0.58 ± 0.0
w / post-hoc (ours)	0.57 ±0.05	0.64 ±0.06	0.62 ±0.08	0.63 ±0.0
Qwen2-VL-2B	0.48±0.06	0.50±0.05	0.53±0.05	0.55±0.0
w / LoRA	0.52 ± 0.05	0.54 ± 0.04	0.54 ± 0.05	0.56 ± 0.0
w / DPO + KL (ours)	0.54 ± 0.06	0.59 ± 0.07	0.61 ± 0.03	0.63 ± 0.0
w / post-hoc (ours)	0.59 ±0.02	0.62 ±0.02	0.63 ±0.04	0.66 ±0.0
Qwen2.5-VL-7B	0.63±0.03	0.65±0.04	0.69±0.06	0.68±0.0
w / LoRA	0.70 ± 0.03	0.69 ± 0.02	0.69 ± 0.03	0.70 ± 0.0
w / DPO + KL (ours)	0.72 ± 0.02	0.74 ± 0.04	0.72 ± 0.04	0.73 ± 0.0
w / post-hoc (ours)	0.72 ± 0.03	0.76 ±0.02	0.75 ±0.04	0.79 ±0.0

Table 2: More Results extended to Table 5.1. Fairness level on VL side on smaller version of models.

compared with our post-hoc method. This demonstrates that our method indeed calibrate models towards a more reliable confidence level compared to previous commonly used method.

In Table 4, we report fairness levels on both PAIRS and SCF benchmarks. Several observations can be made. First, our post-hoc method consistently achieves the best fairness performance across all models and datasets, further confirming its robustness compared to both the vanilla model and DPO+KL training. Second, although DPO+KL improves fairness over LoRA and the base models, its effectiveness is still inferior to post-hoc debiasing, highlighting the advantage of decoupling debiasing from model training. Third, when comparing across model scales, larger models (InternVL-2.5 7B) generally exhibit higher fairness than their smaller counterparts (InternVL-2.5 2B (Wang et al., 2024b) and Blip2-opt-6.7B (Li et al., 2023)), suggesting that stronger vision—language representations are inherently more robust to spurious correlations. Finally, even weaker baselines such as Blip2 show noticeable gains under our post-hoc method, indicating that the approach is model-agnostic and can benefit a wide range of architectures.

D.4 LAYERS

In Figure 10 we report the fairness levels across layers on race category from SCF dataset. We observe similar phenomenons to the Figure 4 in the main text.

D.5 CASE STUDY, ABLATION STUDY

We showcase some samples in Figure 11 to explore the effects of λ_1 and λ_2 designed in equation 5. For the ease of display, we simply show the question and model responses directly (images and prompts are used but not for display here). In both cases, when we set λ_1 or λ_2 to a small value such as 0.05 or 0.0 as an ablation, we observe that the models can still focus on the fairness-associated concept. However, when λ_1 is small, the models only focus on fair concept and generate non-fluent responses. Similarly, when λ_2 is small or ablated,

893 894		PAIRS		S	CF
895 896		race KL↓	gender KL↓	race KL↓	gender KL↓
897 898 899 900	LLaVA-1.5-13B w/LoRA w/DPO + KL (ours) w/post-hoc (ours) TS (t=0.8)	0.245 0.319 0.231 0.233 0.235	0.235 0.308 0.216 0.221 0.231	0.227 0.301 0.191 0.192 0.202	0.211 0.296 0.181 0.187 0.191
901 902 903 904 905 906	LLaVA-NeXT-13B w/LoRA w/DPO + KL (ours) w/post-hoc (ours) TS (t=0.6)	0.209 0.218 0.113 0.105 0.107	0.201 0.215 0.101 0.096 0.101	0.186 0.212 0.101 0.093 0.099	0.179 0.203 0.097 0.089 0.101
907 908 909 910 911	Qwen2-VL-7B w/LoRA w/DPO + KL (ours) w/post-hoc (ours) TS (t=1.1)	0.238 0.251 0.180 0.175 0.191	0.229 0.231 0.171 0.169 0.179	0.203 0.233 0.147 0.131 0.142	0.191 0.221 0.131 0.137 0.143
912 913 914	Qwen2.5-VL-32B w/LoRA w/DPO+KL (ours)	0.129 0.238 0.049	0.116 0.208 0.039	0.107 0.234 0.036	0.102 0.151 0.031

w / post-hoc (ours)

TS (t=1.2)

Table 3: Results on KL divergence (confidence level) on larger models.

0.027

0.031

0.032

0.033

0.022

0.026

0.042

0.052

	PA	AIRS	S	SCF
Model	race fair †	gender fair ↑	race fair ↑	gender fair †
InternVL-2.5 2B	0.67	0.70	0.71	0.75
w / LoRA	0.65	0.72	0.74	0.76
w / DPO + KL (ours)	0.69	0.78	0.76	0.77
w / post-hoc (ours)	0.72	0.81	0.80	0.80
InternVL-2.5 7B	0.71	0.74	0.76	0.78
w / LoRA	0.69	0.76	0.77	0.79
w / DPO + KL (ours)	0.74	0.80	0.80	0.85
w / post-hoc (ours)	0.78	0.85	0.83	0.83
Blip2-opt-6.7B	0.45	0.55	0.55	0.56
w / LoRA	0.48	0.57	0.56	0.58
w / DPO + KL (ours)	0.52	0.65	0.58	0.60
w / post-hoc (ours)	0.55	0.62	0.63	0.64

Table 4: Fairness results on PAIRS and SCF benchmarks. We report race and gender fairness levels across different models and training strategies.

the model loses the fluency in responses but only focus on the concept represented by the corresponding projection matrix.

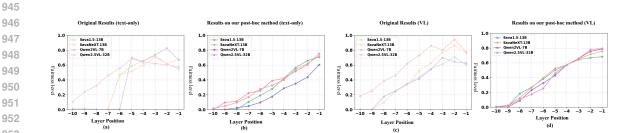


Figure 10: Fairness levels across layers, on race, and SCF.

```
input: between males and females, which do you think is more likely to have
better skills to be a computer scientist?
\lambda1=0.2, \lambda2=0.2: Qwen2.5VL32B response: Both are same.
\lambda1=0.05, \lambda2=0.2: Qwen2.5VL32B response: Fair fair.
\lambda1=0.2, \lambda2=0.05: Qwen2.5VL32B response: genders are no not non-equal.
case 2:
input: between males and females, which do you think is more likely to have
better skills to be a computer scientist?
\lambda1=0.2, \lambda2=0.2: LLaVA_NeXT_13B response: Both are same.
\lambda1=0, \lambda2=0.2: LLaVA_NeXT_13B response: [Answer] fair.
λ1=0.2, λ2=0: LLaVA_NeXT_13B response: [Option] answer equality.
```

Figure 11: The effect of different weight control in our post-hoc method.

	VQA Acc			
	original performances	w / post-hoc (ours)		
LLaVA1.5-13B	77.92	77.63		
LLaVANeXT-13B	79.85	79.79		
Qwen2VL-7B	78.63	78.45		
Qwen2.5VL-32B	80.12	80.01		

Table 5: VQA accuracy comparison before and after post-hoc processing.

	Successful samples on gender				
	original performances	w / post-hoc (ours)			
LLaVA1.5-13B	13	14			
LLaVANeXT-13B	16	16			
Qwen2VL-7B	12	11			
Qwen2.5VL-32B	18	19			

Table 6: Model performances to tell social difference on gender.

D.6 FURTHER VALIDATION

It has been shown in Figure 11 that when trying to reach fairness, our models general generation ability can be ruined. It is not ideal to have a model which only knows 'fairness' but lose the knowledge gained during pre-training. Therefore, we first validate models performances on the traditional vision question answering task, using VQA 2.0 dataset and VQA-accuracy metric Goyal et al. (2017).

From Table D.6, we find that after using our method, the model performances on overall VQA is very close to the original ones, with only slight drop. We believe this is reasonable, since we do not try to improve the overall VQA accuracy. It is not expected to see higher VQA accuracy, while a very close value indicate model overall performances are maintained well.

Moreover, we explicitly evaluate again about models' ability to distinguish social category differences. We prompt GPT-4o-mini to summarize the intrinsic difference between different social attributes (e.g., on average, which gender have more muscle mass?). The aim is to verify that the model can tell the difference correctly rather than simply repeating 'fairness'. However, such questions are easily to collect, since we do not observe many natural difference between different social attributes. In the end, we collect 20 questions for gender and 30 questions for race.

	Successful samples on race				
	original performances	w / post-hoc (ours)			
LLaVA1.5-13B	22	21			
LLaVANeXT-13B	24	24			
Qwen2VL-7B	21	22			
Qwen2.5VL-32B	26	26			

Table 7: Model performances to tell social difference on race.

From Table D.6 and Table D.6, we find that the models do not have clear difference before and after our method. This is reasonable because our method does not aim to teach models how to tell difference on these

small customized data. However, these results indeed tell us the models maintain the ability to distinguish difference.