BENCHMARKING BIAS MITIGATION TOWARD FAIRNESS WITHOUT HARM FROM VISION TO LVLMS

Anonymous authorsPaper under double-blind review

000

001

002003004

010 011

012

013

014

016

017

018

019

021

023

024

025

026

027

028

029

031

033

035

037

038

040

041

042

043

044

046

047

048

051

052

ABSTRACT

Machine learning models trained on real-world data often inherit and amplify biases against certain social groups, raising urgent concerns about their deployment at scale. While numerous bias mitigation methods have been proposed, comparing the effectiveness of bias mitigation methods remains difficult due to heterogeneous datasets, inconsistent fairness metrics, isolated evaluation of vision versus multi-modal models, and insufficient hyperparameter tuning that undermines fair comparisons. We introduce NH-Fair, a unified benchmark for fairness without harm that spans both vision models and large vision-language models (LVLMs) under standardized data, metrics, and training protocols, covering supervised and zero-shot regimes. Our key contributions are: (1) a systematic ERM tuning study that identifies training choices with large influence on both utility and disparities, yielding empirically grounded guidelines to help practitioners reduce expensive hyperparameter tuning space in achieving strong fairness and accuracy; (2) evidence that many debiasing methods do not reliably outperform a well-tuned ERM baseline, whereas a composite data-augmentation method consistently delivers parity gains without sacrificing utility, emerging as a promising practical strategy. (3) an analysis showing that while LVLMs achieve higher average accuracy, they still exhibit subgroup disparities, and gains from scaling are typically smaller than those from architectural or training-protocol choices. NH-Fair provides a reproducible, tuning-aware pipeline for rigorous, harm-aware fairness evaluation. Code: https://anonymous.4open.science/r/submit13349.

1 Introduction

Machine learning (ML) models increasingly shape high-stakes decisions, raising concerns that they replicate or amplify societal biases, leading to unfair outcomes across social groups. Various metrics have been proposed to quantify the model unfairness/bias such as *risk disparity* (Hashimoto et al., 2018), *demographic parity* (Dwork et al., 2012), *equal opportunity*, *equalized odds* (Hardt et al., 2016), *overall accuracy parity* (Berk et al., 2021), and *max-min fairness* (Lahoti et al., 2020), yet each captures different notions and priorities can depend on practical considerations.

Fairness interventions span pre-processing (Qraitem et al., 2023; Sagawa* et al., 2020; Pang et al., 2024; Jang et al., 2021), in-processing (Madras et al., 2018; Xu et al., 2021; Chuang and Mroueh, 2021; Park et al., 2022; Zafar et al., 2019), and post-processing (Hardt et al., 2016; Yin et al., 2024; Dehdashtian et al., 2024; Jung et al., 2024) methods. But they often add complexity, require extensive hyperparameter tuning, or degrade overall accuracy (for example, a model may degrade performance for some groups to meet fairness criteria like demographic or accuracy parity). In safety-critical domains such as healthcare, sacrificing performance for fairness violates ethical principles of beneficence and non-maleficence (Beauchamp and Childress, 1994). Motivated by this, a growing line of work seeks *fairness without harm*, i.e., improving group parity without materially reducing performance for any group (Ustun et al., 2019; Martinez and Bertran, 2019; Yin et al., 2024). Instead of solely enforcing fairness constraints across different groups, these approaches ensure that model performance for every group does not deteriorate.

Despite significant progress, existing approaches were often designed for different problem settings and evaluated in inconsistent, limited environments, making it difficult to assess their general applicability. In particular, several key questions remain unanswered:

- Comparability under a standardized protocol. In fairness research, there are usually inconsistent experimental settings when comparing different methods, e.g., arbitrary choice of optimizer, different hyperparameters, and insufficiently trained baselines. These inconsistencies prevent us from accurately assessing the real utility and fairness performance of state-of-the-art debiasing methods. How do these methods perform in utility, fairness, and overhead, under the same setting with comprehensive hyperparameter sweeping? And can they outperform a carefully trained ERM?
- What makes ERM strong—and how far can it go? There is a practice–research gap: industrial workflows rarely explore the full suite of fairness algorithms and instead prioritize hyperparameter optimization (HPO), like using Ray Tune (Liaw et al., 2018) or Optuna (Akiba et al., 2019), on incumbent models. Yet, exhaustive HPO is computationally expensive. Can we bridge this gap by revealing which training decisions, such as optimizer, model depth, or augmentation, most impact fairness? In addition, if we start from a powerful ERM via principled tuning and selection, can conventional group-disparity mitigations still deliver fairness without materially compromising accuracy? Does the scaling law also apply to AI fairness?
- Where do foundation and multimodal models stand? In the era of foundation models and large pretrained models, are these multi-modal models already fair enough compared with specifically trained models due to pretraining on larger training data with larger model size?

Although some benchmarks have been proposed to make the comparison of previous methods, like MEDFAIR (Zong et al., 2023), FFB (Han et al., 2024), and ABC (Defrance et al., 2024). They often focus on specific domains, omit recent methodological advances, or suffer from limited hyperparameter tuning and dataset diversity. For example, MEDFAIR is restricted to medical datasets and does not evaluate fairness in general vision or multimodal contexts. FFB includes primarily older methods, omitting recent advances in representation learning and data-centric approaches, and lacks sufficient hyperparameter tuning, which may result in suboptimal models. ABCFair focuses on tabular datasets only and uses fixed hyperparameter settings, which limits scalability and may misrepresent method performance. The questions mentioned above are not yet fully addressed. Moreover, existing benchmarks (Xia et al., 2024; Jin et al., 2024a) generally treat classical vision and emerging multimodal models separately, limiting our understanding of how multi-modal pretraining affects fairness relative to other models. We introduce these benchmarks in Appendix A.4.

To bridge this gap, we propose NH-Fair, a comprehensive benchmark for *fairness without harm* in complex image and multimodal settings. NH-Fair unifies evaluation across classical vision models and vision—language models, providing broader insights into fairness comparisons of architectures, pretraining strategies, and model scales. Beyond benchmarking both classical and recent fairness algorithms, we systematically evaluate the role of training choices, the overhead of existing methods, and the performance of state-of-the-art LVLMs. Our goal is to provide not just a benchmark, but actionable insights for developing and deploying fair ML systems. We summarize **contributions and key observations with practical implications for practitioners**:

- 1. ERM and Training Choices Matter. We investigate how training choices affect fairness.
 - Observation. Prior work often overlooks hyperparameter tuning, like optimizer and learning
 rate choice. We show that fixing hyperparameters across methods and datasets can yield unfair
 comparisons. Our findings challenge this practice, where some papers claim "state-of-theart" fairness without sufficient evaluations, and underscore the need for more equitable and
 transparent evaluation protocols in future research.
 - *Takeaway*. Optimizer choice (e.g., SGD, Adam, AdamW, Adagrad) and its learning rate affect both fairness and utility. We recommend focusing on tuning resources here, as these settings have a clear impact. This also applies to selecting the correct pretrained weights, while model depth, batch size, or weight decay have less impact on fairness.
- Revisiting Mitigation Methods. We offer an extensive comparison of recent and classic bias mitigation algorithms, assessing both fairness and accuracy to determine if current methods can achieve equitable performance without compromise.
 - *Observation*. Most fairness-specific algorithms do not significantly outperform a carefully tuned ERM when utilities are accounted for. **Data augmentation** is a simple yet effective strategy that often improves *both* fairness and utility.
 - *Takeaway*. In practice, prioritize augmentation strategies before exploring a wide range of specialized algorithms; this pathway most often achieves fairness without utility loss or overhead.
- 3. LVLMs Are Not Inherently More Fair We extend evaluation to multi-modal models and LVLMs to probe whether large-scale, diverse pretraining reduces disparities or propagates stereotypes.

- Observation. Despite greater data diversity and larger capacity, LVLMs still exhibit fairness issues comparable to task-specific vision models. For example, the Llama series, while achieving strong headline accuracy, remains highly susceptible to dataset-specific biases. In addition, the Model scale shows a weak correlation with fairness improvements.
- Takeaway. Practitioners concerned with fairness should prioritize evaluating different architectures and pretraining choices rather than relying on scaling up models alone.

2 PROBLEM FORMULATION: FAIRNESS WITHOUT HARM

In this paper, we evaluate each ML model based on two criteria: 1) *performance disparity* across different groups, and 2) *negative impact* on model performance caused by fairness interventions. Therefore, we consider the **fairness without harm** problem formulation, which aligns with our goal.

Group fairness notion Consider a dataset of n samples, where sample i is represented as a triple (x_i, y_i, a_i) from the joint distribution P(X, Y, A). Here, $x_i \in \mathbb{R}^d$ is the feature vector, $y_i \in \mathcal{Y}$ is the label, and $a_i \in \mathcal{A}$ is the sensitive attribute (e.g., gender or race). We learn a model $h : \mathbb{R}^d \to \mathcal{Y}$ that achieves strong predictive performance while satisfying a fairness constraint. We evaluate performance using a risk function $R : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$ that quantifies the discrepancy between predictions and labels. To assess disparity across groups, we consider the following fairness criteria.

• Overall Accuracy Parity requires that the classifier's accuracy be equal across groups:

$$\mathbb{P}[h(X) = Y \mid A = a] = \mathbb{P}[h(X) = Y \mid A = a'], \quad \forall a, a' \in \mathcal{A}. \tag{1}$$

We use the accuracy gap between the two groups to denote overall accuracy parity in this paper.

• **Demographic Parity**, also known as statistical parity, requires that the classifier's decisions be independent of the sensitive attribute:

$$\mathbb{P}[h(X) = y \mid A = a] = \mathbb{P}[h(X) = y \mid A = a'], \quad \forall a, a' \in \mathcal{A}, \ y \in \mathcal{Y}. \tag{2}$$

• Equalized Odds requires conditional independence of h(X) and a given the true label, e.g., the true positive rate and the false positive rate are equal across groups:

$$\mathbb{P}[h(X) = y \mid Y = y, A = a] = \mathbb{P}[h(X) = y \mid Y = y, A = a'], \quad \forall a, a' \in \mathcal{A}, \ y \in \mathcal{Y}. \quad (3)$$

• Max-Min Fairness uplifts disadvantaged groups by minimizing the worst group risk:

$$\min_{h \in \mathcal{H}} \max_{a \in \mathcal{A}} \mathbb{E}_{X,Y|A=a} \Big[R\big(h(X), Y\big) \Big]. \tag{4}$$

We do not consider individual fairness (Dwork et al., 2012) or counterfactual fairness (Kusner et al., 2017). The former requires a well-defined similarity function between individuals, which is hard to specify in images, while the latter usually assumes access to causal graphs, which are unavailable in most datasets. Given these constraints, we focus on group fairness.

Fairness without harm Given group fairness notions defined in Sec. 2, enforcing these fairness constraints inevitably harms model performance. Consider Demographic Parity (DP) as an example, if the base rates of outcomes differ across groups ($\mathbb{P}[h(X) = y \mid A = a] \neq \mathbb{P}[h(X) = y \mid A = a']$, $\forall a, a' \in \mathcal{A}$.), enforcing independence between h(X) and A may require distorting predictions to align with group-agnostic rates. Suppose we have two groups and the optimal unconstrained classifier for group 0 satisfies $\mathbb{P}[h(X) = y \mid A = 0] = p_0$. To achieve DP, we must enforce $\mathbb{P}[h(X) = y \mid A = 0] = \mathbb{P}[h(X) = y \mid A = 1] = p$ for some value of p, which could force p to deviate from the group-specific optima p_0 , thereby increasing overall risk. In extreme cases, this might result in a trivial classifier (e.g. h(X) = 1 for all data) that satisfies DP but incurs maximal risk. Thus, fairness interventions risk creating a "race to the bottom," where fairness is not achieved by elevating disadvantaged groups by reducing accuracy for all groups. To avoid this, we adopt the principle of fairness without harm, which augments group fairness with a no-harm condition.Let h_{erm} denote the baseline classifier trained via unconstrained empirical risk minimization (ERM):

$$h_{\text{erm}} = \arg\min_{h \in \mathcal{H}} \sum_{i=1}^{n} R(h(x_i), y_i). \tag{5}$$

The *no-harm* criterion requires that for every group $a \in \mathcal{A}$, the risk incurred by our fairness-enhanced classifier does not exceed that of the baseline:

$$\mathbb{E}_{X,Y|A=a}\Big[R\big(h(X),Y\big)\Big] \le \mathbb{E}_{X,Y|A=a}\Big[R\big(h_{\mathrm{erm}}(X),Y\big)\Big], \quad \forall \, a \in \mathcal{A}. \tag{6}$$

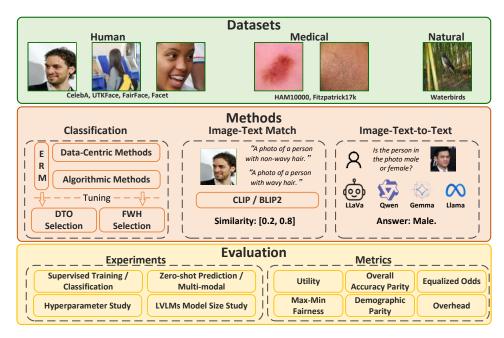


Figure 1: Overview of NH-Fair, evaluating fairness across domains, tasks, models, and methods.

3 NH-FAIR BENCHMARK

3.1 Datasets

We evaluate fairness algorithms on seven publicly available datasets spanning facial attributes, medical imaging, and spurious correlation tests: CelebA, UTKFace, FairFace, Facet, HAM10000, Fitz17k, and Waterbirds. Table 1 summarizes target tasks and sensitive attributes used in our evaluation. Due to the huge computational resource needs, we did not exhaustively use all available sensitive attributes (e.g., gender in UTKFace). Instead, we focused on attributes that exhibited the clear disparity in model predictions, enabling a more effective comparison of existing fairness methods.

Criteria for Dataset Selection. By definition, fairness evaluation requires demographic information (e.g., race, gender, or age) to measure disparities in model performance across subgroups. Therefore, our primary selection criterion is that datasets must include explicit demographic annotations. This ensures that fairness metrics are meaningful and grounded in socially relevant groups. Following this principle, we include the first six datasets. In addition, we considered: application domains, potential sources of bias, such as class imbalance or spurious correlations, and relevance to prior fairness studies. Waterbirds represents a pragmatic exception: it lacks demographic labels but has been widely used in fairness research due to its spurious correlations (e.g., background vs. object) resembling biases. Including it enables consistency with prior work (Reddy et al., 2021; Dehdashtian et al., 2024; Qiang et al., 2024) and provides a test for algorithms in non-demographic bias settings. However, we still suggest not confusing fairness datasets with domain generalization datasets (e.g., Waterbirds, Colored MNIST). While DG datasets probe robustness to distribution shifts, they do not contain socially meaningful sensitive attributes. In contrast, fairness evaluation requires explicit demographic annotations to measure disparities across groups.

Potential Sources of Bias. We use datasets that may introduce bias from multiple sources, including image quality, class imbalance, and spurious correlations. Importantly, each dataset often reflects more than one of these factors simultaneously, making it difficult to attribute disparities to a single cause. For this reason, we do not perform a dataset-level bias source analysis.

3.2 METHODS

We evaluate a diverse set of 12 baseline algorithms, spanning both fairness-specific and general-purpose approaches. To organize them, we group methods into two broad categories: **Data-Centric** and **Algorithmic** methods. **Data-Centric** methods focus on modifying the input distribution

Table 1: Overview of the seven image datasets, detailing the number of samples, classification target, sensitive attribute(s), and approximate imbalance ratios expressed as percentages. In this table, M denotes Male, F denotes Female, W denotes White, B denotes Black, LH denotes Latino Hispanic, EA denotes East Asian, SA denotes Southeast Asian, IN denotes Indian, ME denotes Middle Eastern, N denotes the negative class, and P denotes the positive class.

Dataset	# Samples	Target	Sensitive	Target Ratio	Sensitive Ratio
CelebA (Liu et al., 2015)	200k	Wavy Hair	Gender	68% (F): 32% (M)	41.7% (N): 58.3% (P)
UTKFace (Zhang and Qi, 2017)	23k	Gender	Race	52% (M): 48% (F)	57.6% (W): 42.4% (O)
FairFace (Karkkainen and Joo, 2021)	100k	Ethnicity	Gender	19.1% (W): 14.1% (B): 15.3% (LH) : 14.2% (EA): 12.5% (SA) : 14.2% (IN): 10.7% (ME)	53% (M): 47% (F)
Facet (Gustafson et al., 2023)	30k	Visible Face	Gender	67% (N): 33% (P)	76% (M): 24% (F)
HAM10000 (Maron et al., 2019)	10k	Malignant	Age	86% (N): 14% (P)	71% (Young): 29% (Old)
Fitz17k (Groh et al., 2021)	17k	Malignant	Skin Type	86.5% (N): 13.5% (P)	48% (Light): 52% (Dark)
Waterbirds (Sagawa* et al., 2020)	11k	Bird Species	Background	Train: 76.8% (N): 23.2% (P) Test: 77.8% (Water): 22.2% (Land)	Train: 74.1% (N): 25.9% (P) Test: 50% (Water): 50% (Land)

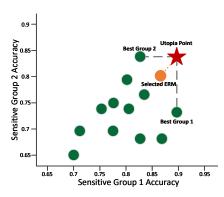
through data augmentation and sampling strategies, including RandAugment (Cubuk et al., 2020), Mixup (Zhang et al., 2018), Resampling (Buda et al., 2018; Sagawa* et al., 2020), Bias Mimicking (BM) (Qraitem et al., 2023), and FIS (Pang et al., 2024). Algorithmic methods, by contrast, alter the training process through adversarial training or fairness-aware objectives. These include Decoupled Classifier (Ustun et al., 2019; Wang et al., 2020), LAFTR (Madras et al., 2018), FSCL (Park et al., 2022), GapReg (Chuang and Mroueh, 2021), MCDP (Jin et al., 2024b), GroupDRO (Sagawa* et al., 2020), and DFR (Kirichenko et al., 2023). Together, these methods span both established baselines and recent state-of-the-art techniques and include pre-processing, in-processing, and post-processing methods. Note that we do not consider causal fairness approaches since most of them assume access to causal graphs or intervention variables. However, these assumptions are rarely feasible in our datasets. A more detailed description of each method is provided in Appendix B.3.

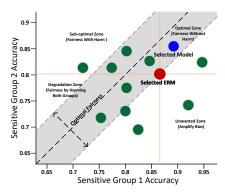
Beyond supervised models, we also evaluate zero-shot predictions from multi-modal models such as CLIP (Radford et al., 2021) and BLIP2 (Li et al., 2023), as well as LVLMs, including LLaVA v1.6 (7B, 13B, 34B) (Liu et al., 2024), Qwen2.5-VL (7B, 32B, 72B) (Bai et al., 2025), Gemma 3 (4B, 12B, 27B) (Team et al., 2025), and Llama models (3.2–11B, 3.2–90B, and 4-Scout–109B) (Grattafiori et al., 2024; Meta, 2025). For CLIP and BLIP2, we construct prompts for text-image matching tasks using pairs of positive and negative labels (e.g., "A photo of a person with non-wavy hair." vs. "A photo of a person with wavy hair."). For LVLMs, we frame image-text-to-text tasks that elicit binary responses, using prompts such as "Is the person in the photo wavy-haired? Answer 'Yes' for wavy hair, 'No' for non-wavy hair." Full prompt templates are provided in Appendix Table 8.

3.3 Model selection

Selecting the best-performing model affects not only overall utility but also disparities across subgroups. Conventional selection methods, which prioritize average accuracy or loss, often reinforce majority-group performance while neglecting minorities. This motivates model selection strategies that explicitly balance utility and fairness. Prior work has explored approaches such as Minimax Pareto Selection (Martinez et al., 2020) and Distance to Optimal (DTO) (Han et al., 2022) to identify fair models. In this work, we adopt a DTO-based strategy to establish a strong ERM baseline, considering both utility and group fairness. We then take the selected ERM model as a baseline and introduce a fairness-without-harm (FWH) procedure to assess whether alternative methods can improve fairness without materially degrading utility compared to the well-established ERM.

- 1) DTO-based selection. The DTO strategy identifies the ERM model that minimizes Euclidean distance to **utopia point**, where each subgroup achieves its best observed performance. This approach aims to find the model that maximizes overall performance while minimizing disparities across subgroups. By establishing this model as the optimal ERM baseline, we provide a reference point for evaluating fairness-aware models. As shown in Figure 2a, the red star denotes the utopia point, where both subgroups reach their respective maximum performance. The model with the smallest DTO value (the shortest distance to the utopia point) is selected as the ERM baseline.
- **2) Fairness-without-harm selection.** Given the optimal ERM, we evaluate other methods by comparing subgroup-specific accuracies. Based on the group accuracy of the selected ERM model, we categorize the models from other methods into four distinct zones:





(a) DTO-based ERM model selection

(b) Fairness-Without-Harm Model Selection

Figure 2: Two-stage model selection. (a) First, each green dot represents a candidate ERM model's performance on two sensitive groups. We select the ERM model (orange circle) whose performance is closest to the utopia point (red star) in Euclidean distance. (b) With the ERM baseline established, we classify models trained by bias-mitigation methods into four zones based on how their subgroup performance compares to the ERM. Starting from the Optimal Zone and moving counterclockwise, we check whether any model is located in the shaded region, which demonstrates improved fairness.

- **Optimal** (Fairness Without Harm): Models here already achieve better utilities for both groups compared with the established baseline ERM. A smaller accuracy gap is desirable to show fairness. Thus, we select the model with the smallest accuracy gap.
- **Sub-optimal** (Fairness by Compromising): Models reduce disparities by decreasing the advantaged group's utility while improving the disadvantaged group. If no candidates exist in **Optimal** zone, we select the model here with the *smallest accuracy gap* to evaluate how it equalizes two groups.
- **Degradation** (Fairness by Harming Both Groups): Models are undesirable since they degrade performance for all groups, achieving fairness at the expense of overall utility. We select a model with the *smallest L2 distance to the optimal ERM* to maintain necessary utility rather than fairness metrics, since the models can significantly decrease utility (random guess) to equalize accuracy.
- **Unwanted**: Models in this zone widen disparities by benefiting the advantaged group and harming the disadvantaged group, which exacerbate unfairness and are not considered.

When evaluating fairness-aware models against the optimal ERM, we follow the selection order: $\mathbf{Optimal} \to \mathbf{Sub\text{-}optimal} \to \mathbf{Degradation}$. This ensures fairness while minimizing performance trade-offs. Using this strategy, we select the model that best meets the fairness-without-harm criteria.

4 Results

Implementation. For all methods, we independently search for the best hyperparameters on each dataset. Specifically, we search over optimizers (SGD and Adam), learning rates, weight decay values, and method-specific hyperparameters. Implementation details are provided in Appendix B.4 and B.5. In total, we spent over 10,000 A100 GPU hours to obtain these benchmarking results.

Metrics. We use **Accuracy** (**ACC**) to evaluate utility on most datasets, while **AUC** is adopted for disease prediction tasks, since medical datasets often exhibit class imbalances. For fairness, we report four metrics: **Overall Accuracy Parity** (**Gap** \downarrow), **Max–Min Fairness** (**Worst** \uparrow), **Demographic Parity** (**DP** \uparrow), and **Equalized Odds** (**EqOdd** \uparrow). They are introduced in Appendix B.2.

4.1 How to tune a fair ERM and why training choices matter

This part is motivated by two observations. First, ML practitioners rarely conduct extensive testing of different fairness-specific algorithms and instead focus on tuning existing models. Yet exhaustive hyperparameter search is computationally expensive, raising the question: which training choices should be prioritized for fairness-sensitive applications? Second, many prior studies report fairness or utility gains while keeping core training parameters fixed across baselines and datasets. But can such practices truly yield fair and reliable comparisons?

To answer these, we conduct a systematic study under different training setups (see Appendix C.1). We find that some choices, most notably initialization (pretraining vs. scratch) and optimizer, consistently impact the fairness—utility trade-off. For example, pretrained models often improve utility without harming disadvantaged groups, and the right optimizer (e.g., SGD on CelebA, Adam on Fitz17k) can shift both utility and subgroup performance. In contrast, factors like batch size, weight decay, or model depth have weaker or inconsistent effects. **These findings suggest that impactful training choices could be prioritized during HPO. In addition, fixing such parameters across baselines and datasets can lead to biased comparisons. It underscores the need for more equitable and transparent evaluation practices in future fairness research.**

4.2 REVISITING MITIGATION METHODS

A well-tuned ERM baseline could match or surpass fairnessspecific algorithms. Table 2 reports utility and fairness outcomes for single-modality models across seven datasets. For each dataset-method pair, results are averaged over five runs, with the best model selected under DTO or FWH criteria. Table 3 extends the comparison to multimodal models. Despite being treated as a "basic" baseline, a well-tuned ERM, selected via DTO, consistently performs competitively across fairness

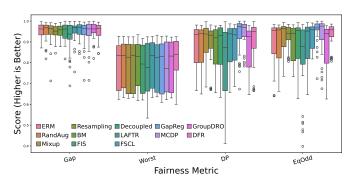


Figure 3: Comparative Analysis of Bias Mitigation Methods.

metrics (Gap, Worst, DP, EqOdd). Figure 3 shows that ERM performs competitively with, or sometimes outperforms, specialized fairness methods under FWH selection, and no single mitigation approach dominates across all datasets.

Table 2: Average utility and fairness results (**standard deviations are reported** in Appendix Table 18). Results better than ERM are highlighted.

Celeba	Dataset	Metric	ERM	RandAug	Mixup	Resampling	BM	FIS	Decoupled	LAFTR	FSCL	GapReg	MCDP	GroupDRO	DFR
CelebA Gap (a) (a) (b) (b) (b) (b) (b) (b) (b) (b) (b) (b		ACC	86.57			86.35			86.35	86.55	85.61	85.62	80.26	86.12	86.58
Part		Worst	83.76	83.89	82.74	83.44	82.86	79.33	83.46	83.67	82.56		77.13	83.50	83.78
$ \begin{array}{c c c c c c c c c c c c c c c c c c c $	ColobA	Gap													
	CelebA														
UTKFace Worst Gagla 91.78 (2.6) (2.34) 91.55 (2.56) (2.56) (2.47) (2.48) 1.97 (2.48) (1.97) (2.10) 2.10 (1.91) (1.91) (2.00) (2.44) (2.63) 2.42 (2.63) (2.74) (2.48) (2.63) 1.97 (3.23) (3.74) (3.		DP	67.20	67.37	70.83	67.39	66.91	66.84	66.68	66.90	69.72	75.91	93.11	66.41	67.30
$ \begin{array}{c c c c c c c c c c c c c c c c c c c $															
FairFace EqOld 97.62 97.61 97.49 97.39 97.51 97.43 97.44 97.44 98.10 98.04 97.06 97.39 97.39 97.39 97.51 97.43 97.44 97.44 98.10 98.04 97.06 97.39 94.78 94.83 94.83 94.51 95.30 95.80 94.78 94.83 94.83 94.51 95.30 95.80 94.78 94.83 94.83 94.85 9															
FairFace	UTKFace														
ACC 66.76 68.37 65.40 65.40 65.66 65.31 67.03 66.44 65.42 65.02 66.06 65.51 63.20	CINFACE														
FairFace Worst Gap Os		DP	94.55	94.83	94.51	95.34	94.34	94.69	96.27	95.44	94.18	95.30	95.80	94.78	94.83
FairFace Gap EqOdd 0.87 0.44 0.93 0.90 0.97 0.53 0.87 0.58 0.87 0.90 0.60 0.59 0.59 0.90 0.60 0.59 0.55 0.50 0.5															
FairFace EqOld 96.22 96.14 96.55 96.83 95.70 95.88 95.14 94.73 97.05 96.15 97.85 96.31 95.81															
Hamilton	FairFace														
ACC 67.55 67.83 67.86 67.56 65.87 67.60 67.33 70.74 67.79 67.01 67.91 67.20 66.87	ranrace														
Facet Warst Gap August Gap		DP	97.61	97.55	97.62	97.80	97.30	97.40	96.88	96.86	98.10	97.51	98.50	97.47	97.43
Facet Gap box															
Facet EqOdd DP 96.40 96.68 ps.75 94.37 ps.30 96.32 ps.09 98.10 ps.09 88.04 ps.09 96.31 ps.09 96.32 ps.09 97.84 ps.09 88.04 ps.09 96.32 ps.09 98.23 ps.09 98.23 ps.09 98.23 ps.09 98.24 ps.09 98.24 ps.09 88.04 ps.09 88.01 ps.09 86.51 ps.09 87.85 ps.09 86.51 ps.09 87.85 ps.09 89.43 ps.09 89.03 ps.09 89.63 ps.09 89.63 ps.09 89.63 ps															
HAM10000 A	Facet														
HAM10000 Ham10000 Ham20000 Ham2000000000000000000000000000000000000	racci														
HAM10000		DP	95.40	95.91	96.71	93.96	95.09	97.84	87.30	95.00	95.66	98.73	98.46	93.16	96.09
$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$															
HAM10000															
Fitz17k AUC 89.74 91.29 90.62 90.76 91.02 88.34 89.63 91.50 99.64 88.55 99.59 91.65 99.72 89.99	HAM10000														
HitzlTik AUC 89.74 91.29 90.62 90.76 91.02 88.34 89.63 91.50 90.83 89.59 91.65 90.72 89.99 Worst 88.39 90.15 89.38 88.99 89.93 87.02 88.45 90.59 89.64 88.52 90.49 90.06 88.57 Gap 2.92 2.51 2.43 3.62 2.34 3.06 2.55 2.95 3.84 1.84 2.87 1.92 2.93 DP 94.46 94.53 94.51 93.31 93.66 95.60 94.09 93.68 97.45 95.47 95.68 94.36 94.30 DP 94.46 94.53 94.51 93.31 93.66 95.60 94.06 94.46 95.32 96.42 96.14 95.24 95.26 ACC 85.63 86.09 87.67 87.35 88.20 83.72 74.64 85.72 86.83 86.45 85.98 85.46 89.83 Waterbirds Gap 2.87 3.14 3.36 4.98 4.48 2.09 20.38 3.56 1.10 1.47 2.31 2.02 1.47 EqOdd 66.53 68.99 81.42 90.87 77.21 65.89 47.31 68.22 90.00 87.48 72.97 67.31 97.79 77.7	11/1/11/0000														
Fitz17k Worst 6ap 88.39 90.15 89.38 88.99 89.93 87.02 88.45 90.59 89.64 88.52 90.49 90.06 88.57 Fitz17k EqOdd 9.29 2.51 2.43 3.62 2.34 3.06 2.55 2.95 3.84 1.84 2.87 1.92 2.93 2.93 2.93 8.84 1.84 2.87 1.92 2.93 2.93 8.62 95.47 95.47 95.64 96.40 95.47 95.47 95.47 95.60 94.06 94.06 94.66 95.32 96.42 96.14 95.24 95.26 96.20 95.27 94.99 93.66 94.66 95.32 96.42 96.14 95.24 95.26 95.26 96.14 95.24 95.26 95.26 96.14 95.24 95.24 95.26 95.26 95.47 96.42 96.14 95.24 95.26 95.26 96.14 95.24 95.26 95.26 96.14 95.24 95.24 95.26		DP	81.48	84.73	88.54	85.05	78.41	88.58	75.15	91.00	77.64	96.74	99.58	83.86	84.65
Fitz17k Gap EqOd 2.92 2.51 2.43 3.62 2.34 3.06 2.55 2.95 3.84 1.84 2.87 1.92 2.93 Fitz17k Gap EqOd 94.92 95.61 96.20 95.27 95.27 95.09 95.17 94.09 95.17 94.09 93.68 97.45 95.47 95.68 94.36 94.30 DP 94.46 94.53 94.51 93.31 93.66 95.60 94.00 94.06 94.46 95.22 96.42 96.14 95.24 95.24 95.04 95.24 95.00 95.00 95.00 94.00 94.00 94.00 94.00 94.00 94.00 94.00 94.00 94.00 94.00 95.00															
Fitz1/k															
Hegodd 94.92 95.61 96.20 95.21 94.99 95.17 94.09 95.08 97.45 95.47 95.88 94.30 94.30 94.00 94.06 94.06 94.00 94.06 95.00 94.00 95.00 96.42 96.00 94.30 95.40 95.40 94.30 94.30 94.30 94.30 94.30 94.30 94.30 94.30 94.00	Fitz17b														
ACC 85.63 86.09 87.67 87.35 88.20 83.72 74.64 85.72 86.83 86.45 85.98 85.46 89.83 Waterbirds Gap 2.87 3.14 3.36 4.98 4.48 2.09 20.38 3.56 1.10 1.47 2.31 2.02 1.47 EqOdd 66.53 68.99 81.42 90.87 77.21 65.89 47.31 68.22 90.00 87.48 72.97 67.31 97.79	I IIZI/K														
Waterbirds Gold 66.99 84.52 84.98 84.95 84.98 82.99 82.09 82.67 64.45 83.94 86.28 85.72 84.83 84.45 89.09 Waterbirds EgOdd 66.53 68.99 81.42 90.87 77.21 65.89 47.31 68.22 90.00 87.48 72.97 67.31 97.79		DP	94.46	94.53	94.51	93.31	93.66	95.60	94.06	94.46	95.32	96.42	96.14	95.24	95.26
Waterbirds Gap EqOdd 2.87 3.14 3.36 4.98 4.48 2.09 20.38 3.56 1.10 1.47 2.31 2.02 1.47 EqOdd 66.53 68.99 81.42 90.87 77.21 65.89 47.31 68.22 90.00 87.48 72.97 67.31 97.79															
Waterbirds EqOdd 66.53 68.99 81.42 90.87 77.21 65.89 47.31 68.22 90.00 87.48 72.97 67.31 97.79															
EqOdd 66.53 68.99 81.42 90.87 77.21 65.89 47.31 68.22 90.00 87.48 72.97 67.31 97.79	Waterbirds														
DP 77.67 77.25 86.00 90.93 81.78 75.30 52.30 77.47 92.53 91.39 80.70 76.91 98.61	mater birds														
		DP	77.67	77.25	86.00	90.93	81.78	75.30	52.30	77.47	92.53	91.39	80.70	76.91	98.61

Data augmentation is a simple yet effective path to fairness without harm. Our results show that fairness improvements do not inherently require sacrificing accuracy. Prior studies like Dutta et al.

(2020) established the theoretical possibility of achieving both. In the paper, **RandAug**, though not designed for bias mitigation, improves both fairness and accuracy across multiple datasets, empirically demonstrating that fairness improvements do not always require a trade-off in model performance. This suggests that increasing data variability can naturally mitigate biases, offering a lightweight and practical strategy for fairness-aware training without the need for explicit fairness constraints.

The utility-fairness trade-offs are pronounced in regularization-based methods. Both GapReg and MCDP account for explicit fairness constraints in the loss functions, directly penalizing subgroup disparities. This design explains why they consistently achieve strong scores on fairness metrics such as EqOdd and DP across datasets. However, the fairness penalty can pull decision boundaries away from the utility-optimal surface, which often results in lower accuracy and occasionally lower Worst-group accuracy. They highlight the classic fairness—utility trade-off: fairness comes at the expense of overall predictive strength. In contrast, the contrastive learning method FSCL, which encourages representations of the same class to cluster across different groups, delivers strong fairness improvements while keeping accuracy competitive, implying that learning fair and robust representations is also a promising path to fairness without harm.

Rethinking spurious correlation datasets for fairness evaluations. In prior studies (Dehdashtian et al., 2024; Qiang et al., 2024; Reddy et al., 2021), datasets with spurious correlations are often used for fairness evaluation since spurious correlations resemble the types of confounding factors that lead to fairness issues. We noticed that, in Table 2, many methods (e.g., Mixup, Resampling, BM...) can achieve both higher utility and fairness on Waterbirds, but it is more difficult to achieve such a gain on fairness datasets with socially meaningful group disparities. This raises a potential concern: spurious correlation might be easier to resolve than fairness-specific distribution shifts. Fairness involves systematic performance gaps across protected groups, which are typically more subtle and persistent than background—object correlations. Over-reliance on datasets like Waterbirds may therefore underestimate the difficulty of fairness challenges and overstate algorithmic effectiveness. While useful for studying robustness to spurious correlations, such datasets are not ideal for fairness evaluation; fairness studies should still focus on those with clear, socially meaningful groups.

4.3 FAIRNESS IN MULTI-MODAL MODELS

We showed that even the best-tuned vision models leave non-trivial subgroup gaps. A natural hypothesis is that these gaps reflect limited data diversity and model capacity. Here, we evaluate pre-trained multi-modal models, which are trained on massive, diverse datasets and are expected to generalize better. Specifically, we consider two families: (1) image—text matching models such as CLIP, BLIP-2, and two CLIP-based debiasing variants (FairerCLIP Dehdashtian et al. (2024), SFID Jung et al. (2024)); and (2) LVLMs such as LLaVA-1.6, Qwen2.5-VL, Gemma 3, and Llama. All models are evaluated in zero-shot predictions. Fig. 4 visualizes the joint landscape of utility and parity, where *Qwen2.5-VL 72B* lies closest to the outer envelope across datasets, indicating the best overall balance between high accuracy and low disparity among the evaluated LVLMs.

Fairness persists as a challenge. From Table 3, we found that fairness outcomes diverge sharply between image—text matchers and newer large models. BLIP-2 and CLIP still exhibit sizable subgroup gaps on CelebA and Waterbirds. Even though debiasing methods like FairerCLIP and SFID can use the validation set to tune the model, they still don't significantly address the fairness issue. As for LVLMs, despite being trained on massive, heterogeneous corpora, mul-

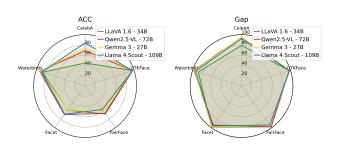


Figure 4: LVLM performance using ACC and 100 - Gap. Other metrics are presented in Figure 8.

timodal models do not universally outperform carefully tuned unimodal baselines in either utility or fairness. On relatively balanced datasets like UTKFace, LVLMs such as Qwen2.5-VL and Gemma-3 deliver strong accuracy while keeping subgroup gaps small. However, on more challenging datasets (CelebA and Facet), they exhibit subgroup disparities, often with worse worst-group accuracy than ERM and sometimes a

larger accuracy gap for certain models. These results emphasize that multi-modal models still inherit and sometimes amplify fairness challenges rather than resolving them.

Scaling is not enough. We also studied different LVLM sizes to evaluate whether the scaling law extends to fairness (Appendix Table 15, simplified in Table 4). Increasing LVLM size improves average accuracy but does not consistently resolve fairness gaps. Larger models (e.g., Gemma-3–27B, Llama3.2-90B) achieve higher average accuracy compared to their smaller counterparts, yet subgroup disparities (Gap) remain non-trivial and in some cases even widen. The fairness gains from scaling are much smaller than those obtained by switching to a different model family, suggesting that the training protocol plays a more decisive role in fairness. Hence, in fairness-sensitive applications, we recommend first exploring model choice before allocating resources to scaling.

Table 3: Comparison of multi-modal models. **Standard deviations** are in Appendix Table 19. We omit HAM10000 and Fitz17k for LVLMs, since they do not yield calibrated probabilities for AUC.

Dataset	Metric	ERM	RandAug	BLIP2	CLIP	FairerCLIP	CLIP-SFID	LLaVA-1.6 34B	Qwen2.5-VL 72B	Gemma3 27B	Llama4-Scout
	ACC	86.57	86.72	47.38	74.07	73.78	72.05	44.83	68.76	74.04	83.71
	Worst	83.76	83.89	36.82	67.43	67.32	66.36	32.69	66.02	72.18	80.54
CelebA	Gap	6.76	6.80	18.09	15.97	15.56	13.69	20.75	4.70	4.47	7.62
	EqOdd	81.91	81.73	97.24	83.72	83.79	95.70	97.41	92.69	92.76	84.81
	DP	67.20	67.37	95.90	81.32	81.06	93.23	91.92	91.71	74.53	71.52
	ACC	92.75	93.19	94.23	96.72	96.79	96.70	97.12	96.78	97.25	97.02
	Worst	91.78	92.19	94.00	95.90	96.05	96.03	96.34	95.76	96.89	96.29
UTKFace	Gap	2.26	2.34	0.45	1.90	1.72	1.55	1.81	2.36	0.85	1.70
	EqOdd	97.62	97.62	99.24	97.96	98.17	98.36	98.16	97.59	99.20	98.27
	DP	94.55	94.83	95.61	93.91	94.27	94.96	95.22	94.33	95.28	94.90
	ACC	66.76	68.37	52.57	57.36	56.81	52.73	66.91	65.57	59.40	56.52
	Worst	66.34	67.69	50.74	57.20	56.41	51.59	66.04	64.73	56.71	53.53
FairFace	Gap	0.87	1.44	3.46	0.34	0.86	2.40	1.84	1.79	5.71	6.35
	EqOdd	96.22	96.14	93.86	97.16	96.78	97.14	96.22	96.26	95.13	96.54
	DP	97.61	97.55	96.89	98.36	98.32	98.46	98.07	97.57	97.16	98.00
	ACC	67.55	67.83	41.16	33.10	33.17	33.26	58.14	68.41	57.75	67.54
	Worst	64.25	64.94	40.40	31.43	31.69	31.54	57.97	63.79	57.25	64.61
Facet	Gap	4.31	3.78	3.22	7.13	6.32	7.37	0.72	6.04	1.42	3.83
	EqOdd	96.47	96.68	93.52	98.72	98.82	99.75	95.90	97.04	96.22	97.55
	DP	95.40	95.91	94.10	98.92	99.00	99.80	93.91	96.15	94.56	96.28
	ACC	85.63	86.09	52.30	78.05	77.77	75.70	86.95	92.41	90.77	88.62
	Worst	84.20	84.52	39.18	74.53	73.49	72.81	81.26	91.27	88.97	85.74
Waterbirds	Gap	2.87	3.14	26.23	7.04	8.56	5.78	11.39	2.28	3.59	5.76
	EqOdd	66.53	68.99	68.24	73.96	73.19	95.57	80.88	94.03	92.99	89.38
	DP	77.67	77.25	63.48	78.12	76.73	94.09	80.46	94.62	93.17	89.44
	AUC	88.35	89.09	38.87	52.15	51.88	52.53	-	-	-	-
	Worst	84.68	84.67	39.00	51.56	52.01	53.41	-	-	-	-
HAM10000	Gap	4.11	4.99	6.89	4.14	3.12	2.24	-	-	-	-
	EqOdd	88.17	88.43	98.19	96.24	95.75	96.19	-	-	_	-
	DP	82.22	84.73	97.89	99.09	98.94	98.32	-	-	-	-
	AUC	89.74	91.29	67.08	69.92	69.81	69.37	-	-	-	-
	Worst	88.39	90.15	66.46	69.78	69.85	69.35	-	-	-	-
Fitz17k	Gap	2.92	2.51	3.74	2.31	2.52	4.73	-	-	-	-
	EqOdd	94.92	95.61	97.06	89.87	87.95	85.88	-	-	-	-
	DP	94.46	94.53	98.40	95.03	93.02	92.19	l –	_	_	_

Table 4: Effect of LVLM model scale on accuracy and fairness metrics.

Dataset Metric	Mataia	Matria LLaVA-1.6			Q	wen2.5-V	L	Gemma 3			Llama		
	Metric	7B	13B	34B	7B	32B	72B	4B	12B	27B	3.2-11B	3.2-90B	4-Scout-109B
	ACC	68.156	69.872	70.79	77.364	72.718	78.386	67.332	73.484	75.842	74.54	76.006	78.682
	Worst	66.104	66.546	66.86	75.65	68.184	76.314	61.58	72.604	74.4	69.322	70.51	76.142
Average	Gap	3.894	5.698	7.302	2.89	7.656	3.434	10.532	2.088	3.208	10.082	11.002	5.052
_	EqOdd	87.912	89.806	93.714	94.98	95.704	95.522	89.086	93.682	95.26	88.01	91.416	93.31
	DP	87.046	88.602	91.916	91.942	96.158	94.876	86.31	90.466	90.94	84.662	87.934	90.028

5 Conclusion

In this study, we introduce NH-Fair, a rigorously curated benchmark for evaluating fairness interventions in image models, and show that the fairness issue remains unresolved. A carefully tuned ERM with the hyperparameter search often rivals specialized debiasing methods. It highlights the crucial role of hyperparameter tuning and model selection in achieving fairness without harm. We further find that utility need not be sacrificed: data augmentation can deliver simultaneous gains in accuracy and subgroup parity. In addition, large vision—language models are not exempt from fairness issues. Their pre-training choices can still imprint spurious correlations that widen gaps. By releasing NH-Fair and all accompanying code, we aim to make fairness results reproducible and to provide the community with a solid baseline on which to build more robust, bias-aware methods.

ETHICS STATEMENT

This work focuses on evaluating fairness in vision and multimodal models. While our study does not involve direct interaction with human subjects, it makes extensive use of publicly available datasets with demographic information. We acknowledge that these datasets may contain label noise, spurious correlations, or demographic imbalances, which themselves reflect broader social biases.

REPRODUCIBILITY STATEMENT

We have taken several steps to ensure the reproducibility of our work. All datasets used in our experiments are publicly available, with detailed descriptions provided in Appendix B.1. Model architectures, training procedures, and hyperparameters are documented in Appendix B.4 and Appendix B.5. The definitions and formulas of all fairness metrics are included in Section 2 and Appendix B.2. We also provide an anonymous code repository linked in the manuscript.

REFERENCES

- Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. In *International Conference on Machine Learning*, pages 1929–1938. PMLR, 2018.
- Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.
- Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 50(1):3–44, 2021.
- Preethi Lahoti, Alex Beutel, Jilin Chen, Kang Lee, Flavien Prost, Nithum Thain, Xuezhi Wang, and Ed Chi. Fairness without demographics through adversarially reweighted learning. *Advances in neural information processing systems*, 33:728–740, 2020.
- Maan Qraitem, Kate Saenko, and Bryan A Plummer. Bias mimicking: A simple sampling approach for bias mitigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20311–20320, 2023.
- Shiori Sagawa*, Pang Wei Koh*, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=ryxGuJrFvS.
- Jinlong Pang, Jialu Wang, Zhaowei Zhu, Yuanshun Yao, Chen Qian, and Yang Liu. Fairness without harm: An influence-guided active sampling approach. *Advances in Neural Information Processing Systems*, 37:61513–61548, 2024.
- Taeuk Jang, Feng Zheng, and Xiaoqian Wang. Constructing a fair classifier with generated fair data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7908–7916, 2021.
- David Madras, Elliot Creager, Toniann Pitassi, and Richard Zemel. Learning adversarially fair and transferable representations. In *International Conference on Machine Learning*, pages 3384–3393. PMLR, 2018.
- Han Xu, Xiaorui Liu, Yaxin Li, Anil Jain, and Jiliang Tang. To be robust or to be fair: Towards fairness in adversarial training. In *International conference on machine learning*, pages 11492–11501. PMLR, 2021.
- Ching-Yao Chuang and Youssef Mroueh. Fair mixup: Fairness via interpolation. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=DN15s5BXeBn.

- Sungho Park, Jewook Lee, Pilhyeon Lee, Sunhee Hwang, Dohyung Kim, and Hyeran Byun. Fair
 contrastive learning for facial attribute classification. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 10389–10398, 2022.
 - Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P Gummadi. Fairness constraints: A flexible approach for fair classification. *The Journal of Machine Learning Research*, 20(1):2737–2778, 2019.
 - Tongxin Yin, Jean-Francois Ton, Ruocheng Guo, Yuanshun Yao, Mingyan Liu, and Yang Liu. Fair classifiers that abstain without harm. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=jvveGAbkVx.
 - Sepehr Dehdashtian, Lan Wang, and Vishnu Boddeti. FairerCLIP: Debiasing CLIP's zero-shot predictions using functions in RKHSs. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=HXoq9EqR9e.
 - 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.
 - Tom L Beauchamp and James F Childress. Principles of biomedical ethics. Edicoes Loyola, 1994.
 - Berk Ustun, Yang Liu, and David Parkes. Fairness without harm: Decoupled classifiers with preference guarantees. In *International Conference on Machine Learning*, pages 6373–6382. PMLR, 2019.
 - N Martinez and M Bertran. Fairness with minimal harm: A pareto-optimal approach for healthcare. *NeurIPS ML4H: Machine Learning for Health*, 2019.
 - Richard Liaw, Eric Liang, Robert Nishihara, Philipp Moritz, Joseph E Gonzalez, and Ion Stoica. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118*, 2018.
 - Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.
 - Yongshuo Zong, Yongxin Yang, and Timothy Hospedales. Medfair: Benchmarking fairness for medical imaging. In *International Conference on Learning Representations (ICLR)*, 2023.
 - Xiaotian Han, Jianfeng Chi, Yu Chen, Qifan Wang, Han Zhao, Na Zou, and Xia Hu. FFB: A fair fairness benchmark for in-processing group fairness methods. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=TzAJbTClAz.
 - MaryBeth Defrance, Maarten Buyl, and Tijl De Bie. ABCFair: an adaptable benchmark approach for comparing fairness methods. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2024. URL https://openreview.net/forum?id=ByknnPI5Km.
 - Peng Xia, Ze Chen, Juanxi Tian, Yangrui Gong, Ruibo Hou, Yue Xu, Zhenbang Wu, Zhiyuan Fan, Yiyang Zhou, Kangyu Zhu, et al. Cares: A comprehensive benchmark of trustworthiness in medical vision language models. *Advances in Neural Information Processing Systems*, 37:140334–140365, 2024.
 - Ruinan Jin, Zikang Xu, Yuan Zhong, Qingsong Yao, DOU QI, S Kevin Zhou, and Xiaoxiao Li. Fairmedfm: fairness benchmarking for medical imaging foundation models. *Advances in Neural Information Processing Systems*, 37:111318–111357, 2024a.
 - Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. *Advances in neural information processing systems*, 30, 2017.

- Charan Reddy, Deepak Sharma, Soroush Mehri, Adriana Romero-Soriano, Samira Shabanian, and Sina Honari. Benchmarking bias mitigation algorithms in representation learning through fairness metrics. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
 - Yao Qiang, Chengyin Li, Prashant Khanduri, and Dongxiao Zhu. Fairness-aware vision transformer via debiased self-attention. In *European Conference on Computer Vision*, pages 358–376. Springer, 2024.
 - Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
 - Song Yang Zhang, Zhifei and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017.
 - 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*, pages 1548–1558, 2021.
 - Laura Gustafson, Chloe Rolland, Nikhila Ravi, Quentin Duval, Aaron Adcock, Cheng-Yang Fu, Melissa Hall, and Candace Ross. Facet: Fairness in computer vision evaluation benchmark. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 20370–20382, 2023.
 - Roman C Maron, Michael Weichenthal, Jochen S Utikal, Achim Hekler, Carola Berking, Axel Hauschild, Alexander H Enk, Sebastian Haferkamp, Joachim Klode, Dirk Schadendorf, et al. Systematic outperformance of 112 dermatologists in multiclass skin cancer image classification by convolutional neural networks. *European Journal of Cancer*, 119:57–65, 2019.
 - Matthew Groh, Caleb Harris, Luis Soenksen, Felix Lau, Rachel Han, Aerin Kim, Arash Koochek, and Omar Badri. Evaluating deep neural networks trained on clinical images in dermatology with the fitzpatrick 17k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1820–1828, 2021.
 - Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: practical automated data augmentation with a reduced search space. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pages 18613–18624, 2020.
 - Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=r1Ddp1-Rb.
 - Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural networks*, 106:249–259, 2018.
 - Zeyu Wang, Klint Qinami, Ioannis Christos Karakozis, Kyle Genova, Prem Nair, Kenji Hata, and Olga Russakovsky. Towards fairness in visual recognition: Effective strategies for bias mitigation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8919–8928, 2020.
 - Jinqiu Jin, Haoxuan Li, and Fuli Feng. On the maximal local disparity of fairness-aware classifiers. In *International Conference on Machine Learning*. PMLR, 2024b.
 - Polina Kirichenko, Pavel Izmailov, and Andrew Gordon Wilson. Last layer re-training is sufficient for robustness to spurious correlations. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=Zb6c8A-Fghk.
 - Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021.

- 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*, pages 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*, pages 26296–26306, 2024.
 - 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. *arXiv* preprint arXiv:2502.13923, 2025.
 - Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.
 - Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
 - AI Meta. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation, april 2025, 2025.
 - Natalia Martinez, Martin Bertran, and Guillermo Sapiro. Minimax pareto fairness: A multi objective perspective. In *International conference on machine learning*, pages 6755–6764. PMLR, 2020.
 - Xudong Han, Timothy Baldwin, and Trevor Cohn. Balancing out bias: Achieving fairness through balanced training. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11335–11350, 2022.
 - Sanghamitra Dutta, Dennis Wei, Hazar Yueksel, Pin-Yu Chen, Sijia Liu, and Kush Varshney. Is there a trade-off between fairness and accuracy? a perspective using mismatched hypothesis testing. In *International conference on machine learning*, pages 2803–2813. PMLR, 2020.
 - Alexandra Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5(2):153–163, 2017.
 - Mohammad Mahdi Khalili, Xueru Zhang, and Mahed Abroshan. Loss balancing for fair supervised learning. In *International Conference on Machine Learning*, pages 16271–16290. PMLR, 2023.
 - Xueru Zhang, Mohammadmahdi Khaliligarekani, Cem Tekin, and Mingyan Liu. Group retention when using machine learning in sequential decision making: the interplay between user dynamics and fairness. *Advances in neural information processing systems*, 32, 2019.
 - Muhammad Bilal Zafar, Isabel Valera, Manuel Rodriguez, Krishna Gummadi, and Adrian Weller. From parity to preference-based notions of fairness in classification. *Advances in neural information processing systems*, 30, 2017.
 - Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems*, 33(1):1–33, 2012.
 - L Elisa Celis, Vijay Keswani, and Nisheeth Vishnoi. Data preprocessing to mitigate bias: A maximum entropy based approach. In *International conference on machine learning*, pages 1349–1359. PMLR, 2020.
 - Aili Shen, Xudong Han, Trevor Cohn, Timothy Baldwin, and Lea Frermann. Contrastive learning for fair representations. *arXiv preprint arXiv:2109.10645*, 2021.
 - Ruijia Wang, Xiao Wang, Chuan Shi, and Le Song. Uncovering the structural fairness in graph contrastive learning. *Advances in neural information processing systems*, 35:32465–32473, 2022.

- Elliot Creager, David Madras, Jörn-Henrik Jacobsen, Marissa Weis, Kevin Swersky, Toniann Pitassi, and Richard Zemel. Flexibly fair representation learning by disentanglement. In *International conference on machine learning*, pages 1436–1445. PMLR, 2019.
- Sungho Park, Sunhee Hwang, Dohyung Kim, and Hyeran Byun. Learning disentangled representation for fair facial attribute classification via fairness-aware information alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2403–2411, 2021.
- Jungsoo Lee, Eungyeup Kim, Juyoung Lee, Jihyeon Lee, and Jaegul Choo. Learning debiased representation via disentangled feature augmentation. Advances in Neural Information Processing Systems, 34:25123–25133, 2021.
- Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. *Advances in neural information processing systems*, 30, 2017.
- Aditya Krishna Menon and Robert C Williamson. The cost of fairness in binary classification. In *Conference on Fairness, accountability and transparency*, pages 107–118. PMLR, 2018.
- Yan Luo, Min Shi, Muhammad Osama Khan, Muhammad Muneeb Afzal, Hao Huang, Shuaihang Yuan, Yu Tian, Luo Song, Ava Kouhana, Tobias Elze, et al. Fairclip: Harnessing fairness in vision-language learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12289–12301, 2024.
- 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*, pages 6820–6829, 2023.
- Walter Gerych, Haoran Zhang, Kimia Hamidieh, Eileen Pan, Maanas K Sharma, Tom Hartvigsen, and Marzyeh Ghassemi. Bendvlm: Test-time debiasing of vision-language embeddings. *Advances in Neural Information Processing Systems*, 37:62480–62502, 2024.
- Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mojsilović, et al. Ai fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5):4–1, 2019.
- Sarah Bird, Miro Dudík, Richard Edgar, Brandon Horn, Roman Lutz, Vanessa Milan, Mehrnoosh Sameki, Hanna Wallach, and Kathleen Walker. Fairlearn: A toolkit for assessing and improving fairness in ai. 2020.
- Yisong Xiao, Xianglong Liu, QianJia Cheng, Zhenfei Yin, Siyuan Liang, Jiapeng Li, Jing Shao, Aishan Liu, and Dacheng Tao. Genderbias-vl: Benchmarking gender bias in vision language models via counterfactual probing: Y. xiao et al. *International Journal of Computer Vision*, pages 1–24, 2025.
- Sibo Wang, Xiangkui Cao, Jie Zhang, Zheng Yuan, Shiguang Shan, Xilin Chen, and Wen Gao. Vlbiasbench: A comprehensive benchmark for evaluating bias in large vision-language model. arXiv preprint arXiv:2406.14194, 2024.
- Leander Girrbach, Stephan Alaniz, Yiran Huang, Trevor Darrell, and Zeynep Akata. Revealing and reducing gender biases in vision and language assistants (VLAs). In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=oStNAMWELS.
- Christophe Denis, Romuald Elie, Mohamed Hebiri, and François Hu. Fairness guarantees in multiclass classification with demographic parity. *Journal of Machine Learning Research*, 25(130): 1–46, 2024.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=YicbFdNTTy.

- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009.
 Léon Bottou. Stochastic gradient descent tricks. In Neural networks: tricks of the trade: second
 - edition, pages 421–436. Springer, 2012.
 - Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2015. URL https://arxiv.org/abs/1412.6980.
 - Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.
 - Rachel Ward, Xiaoxia Wu, and Leon Bottou. Adagrad stepsizes: Sharp convergence over nonconvex landscapes. *Journal of Machine Learning Research*, 21(219):1–30, 2020.

APPENDIX A Related Work Experiment setup B.2 B.3 Methods B.4 B.5 B.6 C Additional experiments C.1.2C.1.3C.1.4C.1.5C.1.6C.3.1Adding new datasets and algorithms to the framework **Supplementary results** Limitations and future directions **Broader impacts** H LLM usage

A RELATED WORK

A.1 FAIRNESS NOTIONS

Various notions have been proposed in the ML literature to measure the unfairness of ML outcomes; they can be roughly classified into the following categories:

- Unawareness: it prohibits the use of sensitive attributes in the training and decision-making under the principle that excluding such features avoids direct discrimination (Dwork et al., 2012).
- Parity-based fairness: it requires certain statistical measures to be equalized across different groups. Prominent examples include Demographic Parity (predicted positive rates should be similar across groups) (Dwork et al., 2012), Equal Opportunity (true positive rates are aligned) (Hardt et al., 2016), Equalized Odds (both false positive and false negative rates are aligned) (Hardt et al., 2016), Predictive Parity (predictive value measures are balanced) (Chouldechova, 2017), Accuracy Parity (overall accuracy remains comparable across groups) (Khalili et al., 2023; Zhang et al., 2019), etc.
- **Preference-based fairness**: inspired by the fair-division and envy-freeness literature in economics, it ensures that given the choice between various sets of decision outcomes, every group of users would collectively prefer its perceived outcomes, regardless of the (dis)parity compared to the other groups (Zafar et al., 2017; Ustun et al., 2019).
- Counterfactual fairness: this notion leverages tools from causal inference and structural causal models to define fairness at the level of individual causal pathways. Intuitively, a model is counterfactually fair if for any individual, the predicted outcome remains unchanged in a "counterfactual world" where the individual belonged to a different demographic group, while all other non-sensitive attributes and causal mechanisms remain fixed (Kusner et al., 2017).
- **Individual fairness**: unlike other notions that ensure fairness at the group level, individual fairness attains fairness at the individual level, by ensuring similar individuals are treated similarly (Dwork et al., 2012).

Our work is mostly related to parity-based fairness. Unlike most existing methods that achieve fairness at the cost of reducing accuracy, we aim to achieve fairness without harm. By assigning group-specific models to different groups, our goal is to reduce the accuracy gap between different groups without sacrificing accuracy for any group (compared to the baseline classifier trained on data collected from all groups).

A.2 APPROACHES TO MITIGATING UNFAIRNESS

A large body of research has focused on mitigating bias in ML models, which can be broadly categorized into three strategies: pre-processing, in-processing, and post-processing interventions.

- **Pre-processing methods** aim to reduce unfairness at the data level. Common approaches include re-weighting or re-sampling training examples to balance demographic groups (Kamiran and Calders, 2012; Qraitem et al., 2023; Sagawa* et al., 2020; Pang et al., 2024), generating synthetic data to fill representation gaps (Jang et al., 2021), or learning fair representations through data transformations (Celis et al., 2020; Chuang and Mroueh, 2021).
- In-processing methods modify the training procedure by directly incorporating fairness constraints. Representative approaches include adversarial training, which encourages representations to be invariant to sensitive attributes (Madras et al., 2018; Xu et al., 2021; Jin et al., 2024b), and fairness-based regularization terms in the loss function (Chuang and Mroueh, 2021; Park et al., 2022; Zafar et al., 2019). More recent work has leveraged contrastive learning (Shen et al., 2021; Park et al., 2022; Wang et al., 2022) or disentangled representations (Creager et al., 2019; Park et al., 2021; Lee et al., 2021) to de-bias learned features.
- Post-processing methods adjust model outputs to better satisfy fairness criteria without retraining. Examples include modifying decision thresholds (Hardt et al., 2016), calibrating prediction scores (Pleiss et al., 2017), or applying confidence-based and surrogate adjustments (Menon and Williamson, 2018; Yin et al., 2024).

Recently, fairness in multi-modal learning has gained increasing attention, especially due to the rising use of vision-and-language models. Several approaches (Dehdashtian et al., 2024; Jung et al.,

2024) have also been proposed to mitigate bias in multi-modal settings. For example, FairCLIP (Luo et al., 2024) minimizes Sinkhorn distance between the two distributions to debias pre-trained vision-language models. (Seth et al., 2023) ensures fair representation from learning additive residual image representations. (Gerych et al., 2024) equalizes the distances between the debiased embeddings and images to achieve test-time fairness for VLMs.

In this paper, we mainly employ pre-processing and in-processing methods (with one post-processing technique Kirichenko et al. (2023)). For clarity of presentation, we therefore re-group these strategies into two broader categories: data-centric methods, which intervene at the dataset or input level, and algorithmic methods, which modify the training process or outputs. This framing better reflects the scope of our study and emphasizes the practical trade-offs practitioners face when choosing between data-level and model-level interventions.

A.3 FAIRNESS WITHOUT HARM

Achieving fairness at the cost of lowering the performance of any group may be undesirable and even prohibited in certain applications. For example, in safety-critical domains such as healthcare, sacrificing model accuracy for fairness is undesirable, as it violates both the beneficence (i.e., doing what is best for patients) and non-maleficence (i.e., avoiding harm) principles in healthcare ethics (Beauchamp and Childress, 1994). Given this concern, some works aim to achieve fairness without negatively impacting model accuracy (Dutta et al., 2020; Ustun et al., 2019; Yin et al., 2024; Pang et al., 2024). Instead of solely enforcing fairness constraints across different groups, these approaches ensure that model performance for every group does not deteriorate. For example, Martinez and Bertran (2019) seeks a Pareto-optimal fair ML model that minimizes performance gaps among groups while preventing *unnecessary* harm (i.e., minimal accuracy reduction for any group). Ustun et al. (2019) leverages individuals' sensitive attributes to train *decoupled* ML models, ensuring each group receives the best possible performance from its assigned model compared to a pooled model (trained on data from all groups) or the decoupled models of other groups. Yin et al. (2024) proposes a method using *abstention*, where a pre-trained ML model selectively defers certain predictions to human decision-makers, thus achieving group fairness without reducing accuracy.

A.4 FAIRNESS BENCHMARKS

A number of toolkits have been developed to standardize fairness evaluation and mitigation. AI Fairness 360 (Bellamy et al., 2019) offers an extensible library of fairness metrics and bias mitigation algorithms across datasets and tasks. Similarly, Fairlearn (Bird et al., 2020) provides practical tools for assessing and improving fairness in machine learning pipelines, with an emphasis on industry adoption. Beyond toolkits, early benchmarking efforts such as Reddy et al. (2021) compared bias mitigation algorithms in representation learning, highlighting trade-offs across fairness metrics. While these efforts established important foundations, they were largely limited to classical ML or relatively simple datasets and did not extend to complex vision or multi-modal contexts.

Subsequent benchmarks have attempted to unify evaluation but remain limited in scope. MEDFAIR (Zong et al., 2023) targets fairness in medical datasets, addressing sensitive healthcare applications in vision tasks. FFB (Han et al., 2024) primarily evaluates older fairness algorithms before 2021 and omits recent advances in representation learning and data-centric strategies, while also lacking systematic hyperparameter tuning. ABCFair (Defrance et al., 2024) focuses on tabular datasets and adopts fixed hyperparameter settings, which restrict scalability and may misrepresent performance in more complex domains.

With the rise of large vision–language models (LVLMs), newer benchmarks have begun to address multimodal fairness. Xia et al. (2024); Jin et al. (2024a) investigate fairness in medical multimodal foundation models. GenderBias-VL Xiao et al. (2025) and VLBiasBench Wang et al. (2024) explore bias in LVLMs but typically cover smaller models, leaving out the larger LVLMs increasingly deployed in real-world applications. VLAs (Girrbach et al., 2025) is the most recent work,, focusing specifically on gender bias in LVLMs and corresponding mitigation strategies.

In contrast, NH-Fair is designed as a general-purpose fairness-without-harm benchmark evaluating multiple demographic attributes (e.g., gender, race, age) across diverse models, datasets, and mitigation strategies. Its distinct contributions are: (1) unifying evaluation across classical vision

and multimodal models; (2) systematically analyzing the role of training settings; and (3) extending coverage to state-of-the-art LVLMs at deployment-relevant scales, ranging from 4B to 109B parameters.

B EXPERIMENT SETUP

B.1 DATASETS DETAILS

Without additional illustration, we randomly split it into training, validation, and testing sets with a ratio of 8:1:1.

CelebA (**Liu et al., 2015**): The CelebFaces Attributes Dataset, known as CelebA, is an extensive collection featuring over 200,000 images of celebrities, with each image annotated for 40 distinct attributes. In our research, we focus on the "wavy hair" attribute as the classification target while treating the "male" attribute as a sensitive feature.

UTKFace (**Zhang and Qi, 2017**): The UTKFace dataset comprises more than 20,000 facial images, each labeled with information on age, gender, and ethnicity. For our analysis, we simplify the race attribute into two categories: "white" and "non-white as the sensitive attribute and take the "gender" as the classification target.

FairFace (**Karkkainen and Joo, 2021**): The FairFace dataset contains over 100,000 facial images annotated with age, gender, and race, emphasizing balanced demographic representation. We employ "race" prediction as the seven-class classification problem, and use "gender" as the sensitive attribute.

Facet (Gustafson et al., 2023): The Facet dataset includes facial images annotated with Fitzpatrick skin types, age, and gender. We utilize the binary attribute "visible face" as the classification target. Since a single data entry in Facet may have multiple skin type annotations, we use "gender" as our attribute for simplicity. Entries with incomplete gender annotations were removed.

HAM10000 (Maron et al., 2019): The HAM10000 dataset is a large collection of dermatoscopic images used for skin lesion classification. We reclassified its 7 diagnostic categories into two broad labels: "benign" and "malignant" following (Maron et al., 2019). Images missing recorded sensitive attributes were excluded from the dataset. HAM10000 dataset includes two sensitive attributes: age and sex. We binarized the age attribute into two categories: "young" and "old." Individuals aged 0-60 years were classified as "young," while those aged 60 years and above were classified as "old."

Fitzpatrick17k (**Groh et al., 2021**): The Fitzpatrick17k dataset (Fitz17k) comprises dermatological images labeled with Fitzpatrick skin types and diagnostic categories. We reclassify diagnoses into "benign" and "malignant" groups, following (Groh et al., 2021), and use skin type as the sensitive attribute, binarized into "lighter" (I-III) and "darker" (IV-VI). Images missing skin type or diagnostic labels were excluded.

Waterbirds (Sagawa* et al., 2020): The Waterbirds dataset contains images of waterbirds and landbirds superimposed on either water or land backgrounds. We classify bird type (waterbird vs. landbird) as the target, with background (water vs. land) serving as the sensitive attribute. Instances with ambiguous background or species labels were removed. We use the train/validation/test split provided with the dataset.

B.2 FAIRNESS METRICS

• Overall Accuracy Parity (Gap): The accuracy/AUC gap between two sensitive groups.

Max-Min Fairness (Worst): The worst accuracy/AUC across two sensitive groups.
 Demographic Parity (DP): For binary classification, we focus only on the positive classification.

• **Demographic Parity (DP):** For binary classification, we focus only on the positive class. For multi-class classifications such as FairFace, we follow the fairness guarantees outlined in (Denis et al., 2024):

$$1 - \max_{y \in [\mathcal{Y}]} \Big| \mathbb{P} \big[h(X) = y | A = a \big] - \mathbb{P} \big[h(X) = y | A = a' \big] \Big|.$$

Table 5: Download URLs for the datasets.

Dataset	Download URL	License
CelebA (Liu et al., 2015)	http://mmlab. ie.cuhk.edu.hk/ projects/CelebA.html	Non-commercial research only
UTKFace (Zhang and Qi, 2017)	https://susanqq. github.io/UTKFace/	Non-commercial research only
FairFace (Karkkainen and Joo, 2021)	<pre>https://github.com/ joojs/fairface</pre>	CC BY 4.0
Facet (Gustafson et al., 2023)	https://ai.meta. com/datasets/ facet-downloads/	Meta Images Research
HAM10000 (Maron et al., 2019)	https://dataverse. harvard.edu/dataset. xhtml?persistentId= doi:10.7910/DVN/ DBW86T	CC BY-NC 4.0
Fitzpatrick17k (Groh et al., 2021)	https://github. com/mattgroh/ fitzpatrick17k	CC BY-NC-SA 4.0
Waterbirds (Sagawa* et al., 2020)	https://github.com/ kohpangwei/group_DRO	No license specified

Equalized Odds (EqOdd): It is evaluated by ensuring parity in the probability of correct classification:

$$1 - \Big| \mathbb{P}[h(X) = y \mid Y = y, A = a] - \mathbb{P}[h(X) = y \mid Y = y, A = a'] \Big|.$$

We calculate this for all classes and take the average across classes as the final outcome.

B.3 METHODS

• Data-centric methods

- RandAugment (Cubuk et al., 2020): RandAugment (denoted as RandAug in the experiment results) is commonly used in semi-supervised and unsupervised learning. It randomly selects and applies a set of data augmentations—such as rotations, translations, or brightness adjustments—to diversify the training data. In this study, we aim to evaluate whether training a model on more diverse data, without using demographic information, can lead to improved fairness.
- Mixup (Zhang et al., 2018): It is a data augmentation technique that combines two training samples and their corresponding labels to help models generalize better and produce smoother predictions. We apply Mixup by blending data from different groups to evaluate its impact on fairness and performance.
- Resampling (Buda et al., 2018; Sagawa* et al., 2020): It balances the dataset distribution by
 either over-sampling underrepresented groups or under-sampling overrepresented groups. This
 is done by assigning weights to each sample, helping to address class or group imbalances.
- Bias Mimicking (Qraitem et al., 2023): It introduces a sampling-based technique that mitigates bias in models by mimicking biased distributions and adjusting them to improve fairness.
- FIS (Pang et al., 2024): It utilizes an influence-guided active sampling approach to enhance the
 representation of sensitive groups while preserving overall performance.

Algorithmic methods

- Decoupled Classifier (Ustun et al., 2019; Wang et al., 2020): It trains separate classifiers for distinct groups and then aggregates their predictions.
- LAFTR (Madras et al., 2018): It uses an adversarial objective designed to disentangle sensitive information from the learned representations.

- FSCL (Park et al., 2022): Fair Supervised Contrastive Loss (FSCL) integrates fairness considerations into a contrastive loss function, aiming to align group-level features without compromising accuracy.
- GapReg (Chuang and Mroueh, 2021): It embeds fairness metrics, such as Demographic Parity, Equal Opportunity, and Equalized Odds, directly into the training objective to balance both accuracy and fairness.
- MCDP (Jin et al., 2024b): It introduces a fairness metric that captures the maximal local disparity by evaluating cumulative ratio disparities across varying neighborhoods. MCDP uses a regularization approach similar to GapReg to optimize the model.
- GroupDRO (Sagawa* et al., 2020): Distributionally Robust Optimization focuses on protecting the worst-case groups through specialized regularization, thereby improving fairness for minorities.
- DFR (Kirichenko et al., 2023): DFR is a post-processing method that retraining only the final layer can mitigate spurious correlations, providing a lightweight and efficient approach for achieving fairer predictions.

• Multi-modal models

1080

1081

1082

1084

1086

1087

1088

1089

1090

1091

1093

1094

1095

1098

1099

1100

1101

1102

1103

1104

1105

1106

11071108

1109

1110

1111

1113

1114

1115

1116

1117

111811191120

1121

1122

1123

1124

1125

1126

1127

1128

11291130

1131

1132

1133

- CLIP (Radford et al., 2021): CLIP (Contrastive Language–Image Pretraining) learns visual concepts from natural language supervision by jointly training on image–text pairs. It has shown strong zero-shot capabilities on image datasets. In this paper, we use the ViT (Dosovitskiy et al., 2021) as the backbone for CLIP. In addition, we consider two CLIP-based post-training debias methods: FairerCLIP (Dehdashtian et al., 2024) and SFID (Jung et al., 2024). FairerCLIP jointly debiases CLIP's image and text representations in reproducing kernel Hilbert spaces, and SFID uses the selective feature imputation for debiasing.
- BLIP2 (Li et al., 2023): BLIP-2 bridges vision and language models using a lightweight Q-Former and pre-trained frozen language models. We use BLIP2 to generate the image embeddings and text embeddings and calculate the embedding similarities to make classifications.
- LLaVA-1.6 (Liu et al., 2024): LLaVA (Large Language and Vision Assistant) is an autoregressive language model that aligns vision features with a language model using instruction tuning.
- Qwen2.5-VL (Bai et al., 2025): Qwen2.5-VL is the latest model of Qwen vision-language series from Alibaba. It integrates a vision encoder and a language model decoder to process multimodal inputs and achieve comparable performance with GPT-40 and Claude 3.5 Sonnet.
- Gemma 3 (Team et al., 2025): Gemma is a family of lightweight open models released by Google, ranging from 1B to 27B parameters. As the 1B model supports only text inputs, we evaluate the multimodal capabilities of the 4B, 12B, and 27B variants in this paper.
- Llama 3.2 (Grattafiori et al., 2024) and Llama 4 (?): Llama 4 is the latest multi-modal model released by Meta. Due to the computation resource limitation, we only evaluate Llama 4 Scout, a smaller variant compared with Llama 4 Maverick, alongside the Llama-3.2 multimodal series.

We present all used models and their sources in Table 6.

B.4 IMPLEMENTATION DETAILS

The experiments were conducted on a supercomputer cluster, where each node is equipped with two AMD EPYC 7643 processors, four NVIDIA A100 GPUs (80GB memory each), and 921GB of RAM. The code is implemented with Python 3.10.12 and PyTorch 2.5.0. All images are resized to 224×224 pixels, and during training, we apply random horizontal flipping for data augmentation. We use a pre-trained ResNet-18 as the backbone initialization to start the training. In addition, we conduct experiments on pretrained weights and model size in Appendix C.

For RandAug, color-based augmentations were excluded during training, as the sensitive attributes in our study include skin type, and such augmentations could inadvertently alter needed features.

The Fairness Influence Selection (FIS) method was originally designed for an active learning setting, where a portion of the dataset remains unlabeled. In its original formulation, FIS selects necessary unlabeled data using a combination of pseudo-labels and ground-truth labels. To adapt it to our setting, we employ ground-truth labels for influence-guided selection. Specifically, the training set is randomly partitioned into two subsets according to a predefined ratio: one subset is used for

Table 6: Summary of multi-modal models

	Model	#Param	URL
CLIP	CLIP-ViT-B/16	150M	https://github.com/openai/CLIP
BLIP-2	BLIP-2 Base	1B	https://github.com/salesforce/ LAVIS
	LLaVA-v1.6-vicuna-7b-hf	7B	https://huggingface.co/LLaVA-hf/ LLaVA-v1.6-vicuna-7b-hf
LLaVA	LLaVA-v1.6-vicuna-13b-hf	13B	https://huggingface.co/LLaVA-hf/ LLaVA-v1.6-vicuna-13b-hf
	LLaVA-v1.6-34b-hf	34B	https://huggingface.co/LLaVA-hf/ LLaVA-v1.6-34b-hf
Qwen	Qwen2.5-VL-7B-Instruct	7B	https://huggingface.co/Qwen/ Qwen2.5-VL-7B-Instruct
	Qwen2.5-VL-32B-Instruct	32B	https://huggingface.co/Qwen/ Qwen2.5-VL-32B-Instruct
	Qwen2.5-VL-72B-Instruct	72B	https://huggingface.co/Qwen/ Qwen2.5-VL-72B-Instruct
	gemma-3-4b-it	4B	https://huggingface.co/google/ gemma-3-4b-it
Gemma	gemma-3-12b-it	12B	https://huggingface.co/google/ gemma-3-12b-it
	gemma-3-27b-it	27B	https://huggingface.co/google/gemma-3-27b-it
	Llama-3.2-11B-Vision-Instruct	11B	https://huggingface. co/meta-Llama/Llama-3.
Llama	Llama-3.2-90B-Vision-Instruct	90B	2-11B-Vision-Instruct https://huggingface. co/meta-Llama/Llama-3. 2-90B-Vision-Instruct
	Llama 4 Scout (17Bx16E)	109B	https://huggingface. co/meta-Llama/Llama% 204-Scout-17B-16E-Instruct

supervised training, while during validation, the model assesses the influence of data in the second subset and selectively incorporates a subset of these samples into the training set to achieve fairness without harm. The split ratio is treated as a hyperparameter, as it varies across different datasets.

For Fairness-Sensitive Contrastive Learning (FSCL), training was divided into two phases: The first 60 epochs perform contrastive learning. The subsequent 40 epochs were used for classifier training.

For GapReg and MCDP, the differential privacy (DP) loss function was extended to accommodate multi-class classification, particularly for the FairFace dataset. This extension is based on the multi-class demographic parity formulation proposed by (Denis et al., 2024).

For CLIP and BLIP2, we use their public implementations. For LLVMs, we adopt the open-source weights from Huggingface for all models with BF16 or FP16 precision, depending on their suggested model loading.

B.5 Hyperparameters

We perform extensive hyperparameter optimization using Bayesian hyperparameter search via Weights & Biases on the validation results. The default batch size is set to 256 for all experiments. For the SGD optimizer, we deploy a StepLR scheduler and set momentum to 0.9, where the learning rate is reduced by a factor of 0.1 every 30 epochs. The number of hyperparameter search iterations varies with the complexity of each method–dataset combination, subject to a tuning budget of 100-200 searches per method–dataset pair. The search space includes both discrete choices and continuous ranges, as detailed in Table 7. $\mathcal U$ denotes a uniform distribution, and $\log \mathcal U$ denotes a log-uniform distribution. The benchmarking process required a total of approximately 1.1 GPU years. All methods

were trained for 100 epochs with early stopping to prevent overfitting. Early stopping is used if the validation loss or validation accuracy/AUC does not improve after 10 epochs.

Table 7: Hyperparameter search space for all methods.

Category	Hyperparameter	Search Space
Common Hyperparameters	Learning Rate Weight Decay Optimizer	$ \begin{array}{c} \log \mathcal{U}(10^{-5}, 10^{-2}) \\ \{0, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\} \\ \{\text{Adam, SGD}\} \end{array} $
Mixup	Mixup loss coefficient	U(0.1, 10)
Resampling	Resampling method	$\{ \text{Group, Group} \times \text{Class} \}$
BM	Sampling mode	{none, us, uw, os}
FIS	Label ratio Fairness metric	{0.1, 0.3, 0.5} {dp, eop, eod}
LAFTR	Class coefficient Fairness coefficient	$\mathcal{U}(0.1,1) \\ \mathcal{U}(0.1,1)$
FSCL	Group normalization	{0, 1}
GapReg	Fairness regularization objective Regularization coefficient	$ \begin{cases} dp, eop, eod \rbrace \\ \mathcal{U}(0.01, 5) \end{cases} $
MCDP	Regularization coefficient Temperature	<i>U</i> (0.01, 5) {5, 10, 20, 50, 100}
GroupDRO	Alpha Gamma	$\mathcal{U}(0.01, 5) \ \mathcal{U}(0.01, 5)$
DFR	Tune class weights Add train data Regularization C options Class weight options	{0, 1} {0, 1} {L1, L2} {3, 1, 0.7, 0.3, 0.1, 0.07, 0.03, 0.01, 0.003} {1, 2, 3, 10, 100, 300, 1000}
FairerCLIP	$\begin{array}{c} \tau_i \\ \tau_{z,i} \\ \tau_t \\ \tau_{zt} \\ \text{RFF Dim} \end{array}$	$\mathcal{U}(0,1)$ $\mathcal{U}(0,1)$ $\mathcal{U}(0,1)$ $\mathcal{U}(0,1)$ $\mathcal{U}(0,1)$ $\{1000,2000,3000,4000,8000\}$
SFID	Confidence threshold Number of Pruned Feature	$\mathcal{U}(0.5, 1)$ {50,75,100}

B.6 PROMPTS

Table 8 lists the zero-shot prompts used throughout our study. For image—text matching models (CLIP, BLIP-2, FairerCLIP, SFID), we use one template sentence per class and compute the image—caption similarity, then pick the label of the highest-scoring caption as the prediction. For large vision—language models, we feed the prompt with image input to the model and then parse the model's short text answer.

Table 8: Prompts used across datasets for text-image matching and visual-language model (image-text-to-text) inference.

Dataset	Prompts
CelebA	Image-Text Matching: ["A photo of a person with non-wavy hair.", "A photo of a person with wavy hair."] Image-Text-to-Text: Is the person in the photo wavy-haired? Only Answer 'Yes' for wavy hair, 'No' for non-wavy hair.
UTKFace	Image-Text Matching: ["A photo of a male.", "A photo of a female."] Image-Text-to-Text: Is the person in the photo male or female? Only Answer 'Male' or 'Female'.
FairFace	Image-Text Matching: ["A photo of a White person.", "A photo of a Black person.", "A photo of a Latino or Hispanic person.", "A photo of an East Asian person.", "A photo of a Southeast Asian person.", "A photo of an Indian person.", "A photo of a Middle Eastern person."] Image-Text-to-Text: Based on the photo, classify the race of the person into: 0: White, 1: Black, 2: Latino or Hispanic, 3: East Asian, 4: Southeast Asian, 5: Indian, 6: Middle Eastern. Please answer with the corresponding category number only.
Facet	Image-Text Matching: "A photo of a person with non-visible face." "A photo of a person with visible face." Image-Text-to-Text: Does the photo show a visible face? Only answer 'Yes' if the face is visible, otherwise answer 'No'
Waterbirds	Image-Text Matching: "A photo of a landbird." "A photo of a waterbird." Image-Text-to-Text: Is the bird in the photo a landbird or a waterbird? Only answer 'Landbird' or 'Waterbird'.
HAM10000 / Fitz17k	Image-Text Matching: "A photo of a benign skin condition." "A photo of a malignant skin condition."

C ADDITIONAL EXPERIMENTS

C.1 TRAINING CHOICE EXPLORATIONS

In this paper, we argue that a carefully tuned ERM often rivals specialized debiasing methods. We then organize the further studies to mirror the life cycle of a standard vision pipeline. We start with initialization choices from pre-trained weights versus training from scratch, because the features a model begins with strongly constrain downstream bias. Next, we study optimizer choice to challenge prior work that often fixes optimizers across different settings and evaluate their influence on both utility and fairness. We then turn to other common training choices like batch size, weight decay, and model/checkpoint selection. Together, these investigations shed light on fairness outcomes and also help reduce the search space for future hyperparameter optimization, providing principled guidance on which training choices matter most.

C.1.1 IMPACT OF PRETRAINED WEIGHTS

Loading pretrained weights improves overall utility without harming fairness. In this part, we investigate the impact of loading pretrained weights on ImageNet (Deng et al., 2009) before training on fairness and utility on ERM. For clarity, Figure 5 extracts the most representative panels from the full grid in Figure 6. Each panel plots models trained with pretraining (red circles, "Yes") and from scratch (blue circles, "No") across our seven datasets while sweeping the same hyperparameter grid, making the influence of weight initialization easy to see at a glance. The horizontal axis of each plot measures subgroup performance disparity (e.g., "Gap AUC" or "Gap ACC"), while the vertical axis measures overall utility (e.g., "Overall AUC" or "Overall ACC"). In all cases, points closer to the top-left corner indicate both higher overall performance and smaller subgroup gaps, which represent more favorable fairness-utility trade-offs. A general trend emerges showing that models initialized with pretrained weights tend to achieve higher overall performance than those trained from scratch. This benefit is especially apparent in datasets with relatively small training sets, such as Fitz17k and Waterbirds. However, pretrained models do not consistently outperform models trained from scratch in closing the subgroup performance gap. This observation is further supported by other fairness metrics (such as DP, EqOdd, and Worst) reported in the Figure 6, where the full comparison is provided. Despite this mixed impact on fairness, pretrained models improve overall utility without harming any particular subgroup, particularly benefiting disadvantaged groups (as seen in the Worst metric in the appendix). This aligns with our "fairness without harm" principle. As a result, we adopt pretrained models in our benchmark for their superior overall performance while maintaining acceptable fairness levels.

C.1.2 IMPACT OF OPTIMIZER CHOICE ON FAIRNESS AND UTILITY

Using a fixed optimizer across methods and datasets can lead to unfair comparisons. Many prior studies use a single optimizer across datasets and methods, sometimes even without fine-tuning learning rates. Based on this, we conducted this study on two baseline methods and two datasets to provide empirical evidence that using a fixed optimizer across methods and datasets can lead to unfair comparisons. Our findings in Figure 7 indicate that different datasets respond differently to various optimizers, and selecting the right one can improve both utility and fairness. For instance, models trained on CelebA with SGD tend to achieve better utility and smaller accuracy gaps, resulting in a more balanced model compared to those trained with Adam. This is evident in the results, where SGD-based models cluster around higher accuracy values with lower fairness gaps. On the other hand, in the Fitz17k dataset, models trained with Adam perform better, achieving higher utility while maintaining competitive fairness scores. This variation in optimizer performance highlights that a one-size-fits-all approach can lead to suboptimal results, especially in fairness-sensitive settings. Our findings underscore the need for more equitable and transparent evaluation practices in future fairness research.

Further Analysis (ERM Focus). In the main text, due to computational constraints, we primarily experimented with SGD (Bottou, 2012) and Adam (Kingma and Ba, 2015). Here, we extend the study to include two additional optimizers, AdamW (Loshchilov and Hutter, 2019) and Adagrad (Ward et al., 2020), in order to better understand their impact on ERM performance. Across datasets, optimizer choice exhibits a more pronounced and consistent influence on both utility and fairness than model size or batch size. For example, on Waterbirds, AdamW and Adagrad achieve both higher

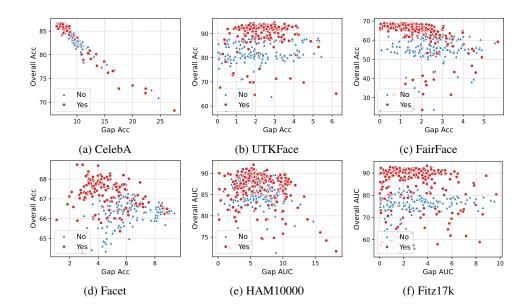


Figure 5: Comparison of utility gap for the pretrained weights. More left indicates greater fairness (smaller accuracy gap), while higher values indicate better performance.

Table 9: Performance comparison across different optimizers.

Dataset	Optimizer	Utility (Acc or AUC)	Worst	Gap	EqOdd	DP
•	SGD	86.57 ± 0.2	83.76 ± 0.2	6.76 ± 0.3	81.91 ± 0.6	67.20 ± 0.7
	Adam	86.11 ± 0.4	82.85 ± 0.6	7.89 ± 0.3	80.46 ± 3.2	67.22 ± 2.8
CelebA	AdamW	86.48 ± 0.1	83.30 ± 0.2	7.67 ± 0.1	79.27 ± 2.4	66.03 ± 1.9
	Adagrad	86.65 ± 0.1	83.62 ± 0.3	7.33 ± 0.3	81.39 ± 0.9	67.90 ± 0.3
	SGD	91.95 ± 0.0	90.51 ± 0.2	3.30 ± 0.6	96.76 ± 0.6	95.56 ± 0.3
	Adam	92.75 ± 0.5	91.78 ± 0.6	2.26 ± 0.6	97.62 ± 0.5	94.55 ± 1.2
UTKFace	AdamW	92.85 ± 0.3	91.67 ± 0.4	2.72 ± 0.7	97.30 ± 0.7	95.70 ± 0.2
	Adagrad	92.84 ± 0.2	91.27 ± 0.3	3.59 ± 0.2	96.45 ± 0.1	95.17 ± 0.4
	SGD	65.68 ± 0.4	64.64 ± 0.5	1.86 ± 0.3	95.98 ± 0.4	97.36 ± 0.2
	Adam	66.76 ± 0.2	66.34 ± 0.4	0.87 ± 0.5	96.22 ± 0.6	97.61 ± 0.4
FairFace	AdamW	67.09 ± 0.4	66.79 ± 0.4	0.64 ± 0.1	96.50 ± 0.4	97.69 ± 0.2
	Adagrad	67.54 ± 0.4	66.83 ± 0.7	1.50 ± 0.6	96.07 ± 0.4	97.65 ± 0.2
	SGD	67.55 ± 0.4	64.25 ± 1.0	4.31 ± 1.1	96.47 ± 1.2	95.40 ± 1.1
	Adam	66.34 ± 0.9	63.09 ± 1.3	4.23 ± 0.9	94.60 ± 1.4	94.17 ± 1.2
Facet	AdamW	67.46 ± 0.2	62.23 ± 0.7	6.80 ± 1.1	95.67 ± 1.2	94.62 ± 0.9
	Adagrad	67.62 ± 0.1	62.99 ± 0.2	6.02 ± 0.3	97.15 ± 0.9	95.95 ± 0.8
	SGD	88.35 ± 1.8	84.68 ± 2.0	4.11 ± 2.1	88.17 ± 3.1	82.22 ± 4.8
	Adam	86.47 ± 1.1	81.74 ± 1.7	4.53 ± 1.7	92.55 ± 2.4	87.69 ± 2.7
HAM10000	AdamW	88.18 ± 0.7	84.38 ± 0.7	3.40 ± 1.1	89.02 ± 6.0	80.23 ± 7.6
	Adagrad	87.00 ± 0.9	81.98 ± 2.9	5.27 ± 3.2	91.33 ± 3.6	85.52 ± 1.0
	SGD	89.74 ± 1.0	88.39 ± 1.1	2.92 ± 1.1	94.92 ± 2.5	94.46 ± 1.4
	Adam	89.62 ± 0.5	87.91 ± 0.5	2.34 ± 0.4	94.90 ± 0.9	91.74 ± 0.8
Fitz17k	AdamW	91.37 ± 0.3	90.45 ± 0.3	2.11 ± 0.9	93.91 ± 1.6	92.32 ± 0.6
	Adagrad	91.43 ± 0.3	90.40 ± 0.6	3.51 ± 1.7	96.09 ± 0.9	93.40 ± 0.5
	SGD	85.45 ± 0.9	83.72 ± 1.7	3.57 ± 1.8	68.23 ± 2.0	76.09 ± 2.0
	Adam	85.63 ± 1.4	84.20 ± 0.9	2.87 ± 0.9	66.53 ± 3.3	77.67 ± 3.5
Waterbirds	AdamW	87.09 ± 0.5	86.25 ± 0.3	1.68 ± 0.9	71.90 ± 0.4	81.54 ± 0.8
	Adagrad	87.77 ± 0.6	87.10 ± 0.9	1.35 ± 1.5	74.00 ± 3.0	81.35 ± 2.0

overall accuracy and smaller subgroup gaps compared to SGD and Adam, indicating a more favorable fairness—utility trade-off. These results suggest that the relative advantages of each optimizer are highly dataset-dependent, yet their effects on utility and fairness are clearer than those of other hyperparameters. Thus, when conducting hyperparameter optimizations, the choice of optimizer could be prioritized for fairness-sensitive evaluations.

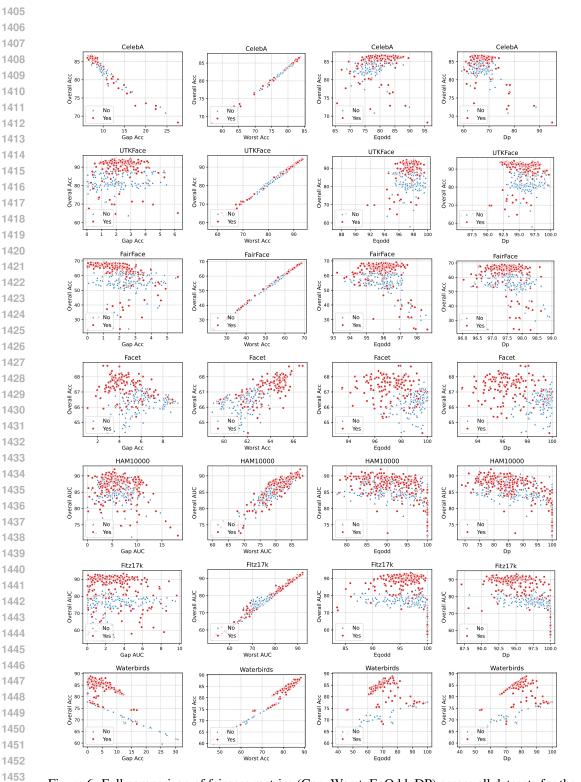


Figure 6: Full comparison of fairness metrics (Gap, Worst, EqOdd, DP) across all datasets for the pretrained weights.

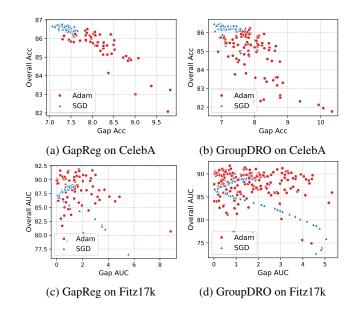


Figure 7: The effect of optimizer choice on fairness and utility trade-offs across CelebA and Fitz17k. The points from the same color are models trained with different learning rates and weight decays.

C.1.3 IMPACT OF BATCH SIZE

Batch size has limited effect on overall utility but can affect fairness slightly. In this benchmark, we use a batch size of 256 for all methods. We also evaluate if batch size have an influence on model fairness. We conduct experiments on seven datasets with different batch sizes ranging from 32 to 1024 and report the results in Table 10. Across all datasets, utility remains relatively stable as batch size increases. However, worst-case performance, which represents accuracy for the most disadvantaged group, fluctuates more significantly. In some cases, such as CelebA and HAM10000, worst-case performance tends to decrease slightly with larger batch sizes. This indicates that larger batch sizes might lead to increased disparities between advantaged and disadvantaged groups. Equalized Odds and Demographic Parity vary highly depending on specific datasets. For datasets like CelebA and Waterbirds, EqOdd tends to decrease as batch size increases, while EqOdd and DP increase with larger batch sizes in HAM10000 and Fitz17k. The results show that batch size can influence fairness, but not in a consistent way. However, since its influence on fairness is relatively minor and inconsistent, batch size can be given lower priority in future hyperparameter searches.

C.1.4 IMPACT OF WEIGHT DECAY

Weight decay is a critical hyperparameter that affects both model generalization and utility. The results in Table 11 suggest that while weight decay influences fairness, its effects vary significantly across different datasets. L2 regularization tends to maintain or slightly improve fairness metrics, whereas L1 regularization can, in some cases, amplify disparities, particularly in datasets with imbalanced subgroup distributions.

C.1.5 IMPACT OF MODEL SIZE

Increasing model size does not consistently improve fairness. We analyze the effect of model size on fairness by comparing different ResNet architectures (ResNet-18, ResNet-34, ResNet-50, and ResNet-101) across UTKFace, FairFace, and HAM10000 datasets in Table 12. We observe that larger models tend to have slightly higher utility (overall accuracy or AUC), however, this improvement is not consistent in the disadvantaged group in terms of the worst acc. For fairness metrics, larger models exhibit varying trends across datasets without an obvious pattern. The results suggest that increasing model size does not consistently improve fairness for single-modality classic supervised learning.

Table 10: Performance comparison across different batch sizes.

Dataset	Batch Size	Utility (Acc or AUC)	Worst	Gap	EqOdd	DP
	32	86.48 ± 0.3	83.59 ± 0.3	6.94 ± 0.4	80.82 ± 1.6	66.77 ± 1.2
	64	86.63 ± 0.2	83.79 ± 0.4	6.82 ± 0.5	81.72 ± 1.9	67.37 ± 1.3
Colob A	128	86.49 ± 0.1	83.69 ± 0.2	6.73 ± 0.6	81.80 ± 0.9	67.23 ± 1.1
CelebA	256	86.57 ± 0.2	83.76 ± 0.2	6.76 ± 0.3	81.91 ± 0.6	67.20 ± 0.7
	512	86.50 ± 0.1	83.60 ± 0.2	6.96 ± 0.4	81.57 ± 0.9	67.09 ± 0.8
	1024	86.52 ± 0.1	83.67 ± 0.1	6.86 ± 0.3	81.60 ± 0.6	67.28 ± 0.7
	32	93.07 ± 0.5	92.15 ± 0.7	2.12 ± 1.0	97.37 ± 0.4	94.21 ± 1.7
	64	93.28 ± 0.2	92.53 ± 0.3	1.74 ± 0.5	98.22 ± 0.5	94.80 ± 2.0
UTKFace	128	91.91 ± 0.0	90.46 ± 0.0	3.40 ± 0.0	96.51 ± 0.0	96.19 ± 0.0
UIKrace	256	92.75 ± 0.5	91.78 ± 0.6	2.26 ± 0.6	97.62 ± 0.5	94.55 ± 1.2
	512	93.30 ± 0.4	92.48 ± 0.8	1.90 ± 0.9	98.10 ± 0.9	95.84 ± 1.2
	1024	92.80 ± 0.5	91.62 ± 0.8	2.75 ± 0.8	97.25 ± 0.8	95.13 ± 1.9
	32	67.49 ± 0.3	66.64 ± 0.4	1.80 ± 0.6	96.39 ± 0.4	97.80 ± 0.3
	64	67.75 ± 0.5	67.43 ± 0.5	0.67 ± 0.5	96.13 ± 0.4	97.24 ± 0.3
FairFace	128	66.76 ± 0.5	66.40 ± 0.7	0.76 ± 0.6	96.32 ± 0.7	97.56 ± 0.4
ranrace	256	66.76 ± 0.2	66.34 ± 0.4	0.87 ± 0.5	96.22 ± 0.6	97.61 ± 0.4
	512	66.63 ± 0.4	65.95 ± 0.6	1.44 ± 0.5	96.15 ± 0.6	97.59 ± 0.3
	1024	65.77 ± 0.7	65.19 ± 0.8	1.23 ± 0.4	96.05 ± 0.5	97.41 ± 0.3
	32	67.55 ± 0.3	64.11 ± 0.6	4.50 ± 0.9	96.33 ± 0.9	95.31 ± 0.9
	64	67.55 ± 0.5	64.23 ± 0.8	4.35 ± 1.1	97.11 ± 1.1	96.18 ± 1.0
Facet	128	67.68 ± 0.1	63.61 ± 0.7	5.31 ± 1.0	97.08 ± 0.7	96.05 ± 0.7
racei	256	67.55 ± 0.4	64.25 ± 1.0	4.31 ± 1.1	96.47 ± 1.2	95.40 ± 1.1
	512	67.58 ± 0.1	63.85 ± 0.7	4.88 ± 1.1	97.31 ± 0.8	97.14 ± 0.9
	1024	67.08 ± 0.3	62.71 ± 0.9	5.72 ± 1.2	98.11 ± 1.1	97.87 ± 1.2
	32	89.35 ± 2.2	85.39 ± 3.8	4.30 ± 2.1	87.04 ± 3.6	82.25 ± 4.3
	64	89.32 ± 1.4	86.08 ± 3.1	3.79 ± 3.0	88.25 ± 4.4	82.91 ± 3.4
HAM10000	128	89.21 ± 1.7	84.70 ± 3.6	5.82 ± 3.0	86.80 ± 4.1	81.80 ± 0.9
11/1/11/0000	256	88.35 ± 1.8	84.68 ± 2.0	4.11 ± 2.1	88.17 ± 3.1	82.22 ± 4.8
	512	88.21 ± 1.7	83.61 ± 2.7	5.08 ± 2.8	89.07 ± 3.9	84.73 ± 3.3
	1024	88.17 ± 1.3	83.30 ± 2.3	5.40 ± 2.9	87.22 ± 5.6	84.91 ± 2.3
	32	90.62 ± 0.8	89.45 ± 0.8	2.80 ± 1.3	96.50 ± 2.7	94.55 ± 1.1
	64	90.36 ± 0.7	89.02 ± 1.3	2.73 ± 1.6	96.13 ± 2.1	94.85 ± 1.0
Fitz17k	128	89.92 ± 0.7	88.50 ± 0.8	3.16 ± 1.2	95.51 ± 2.5	94.54 ± 1.2
1102171	256	89.74 ± 1.0	88.39 ± 1.1	2.92 ± 1.1	94.92 ± 2.5	94.46 ± 1.4
	512	88.94 ± 1.1	87.70 ± 1.2	2.61 ± 1.3	95.65 ± 2.4	95.27 ± 1.2
	1024	86.97 ± 1.4	85.49 ± 1.6	2.82 ± 1.6	95.95 ± 1.1	96.49 ± 0.8
	32	87.29 ± 1.3	85.81 ± 1.1	2.98 ± 0.9	73.88 ± 3.1	82.33 ± 3.0
	64	87.70 ± 0.6	86.36 ± 1.0	2.68 ± 1.4	74.60 ± 1.8	83.17 ± 1.8
Waterbirds	128	88.14 ± 0.6	87.02 ± 0.8	2.24 ± 1.7	75.69 ± 2.0	84.11 ± 1.7
	256	85.63 ± 1.4	84.20 ± 0.9	2.87 ± 0.9	66.53 ± 3.3	77.67 ± 3.5
	512	87.11 ± 1.3	86.03 ± 0.8	2.16 ± 1.4	72.69 ± 3.7	81.45 ± 2.9
	1024	83.68 ± 1.2	81.35 ± 2.4	4.65 ± 2.8	63.98 ± 1.6	73.77 ± 3.2

C.1.6 IMPACT OF MODEL SELECTIONS

We use ERM, RandAug, and GroupDRO as examples to compare the model selection methods in Table 13 since their differences are more obvious to observe compared with others. The DTO-based selection method prioritizes models that maximize utility for all groups, ensuring that no subgroup is disproportionately disadvantaged. On the other hand, the FWH selection strategy emphasizes fairness without significantly harming utility based on the DTO-based method selected ERM results, focusing on reducing disparities without major accuracy degradation. The *M* in the Table denotes using maximal overall utility. FWH selection achieves a smaller fairness gap across multiple datasets. These results show that prioritizing fairness constraints does not necessarily mean sacrificing model performance, as long as careful selection strategies are employed. However, since the selection is performed on the validation set, it still faces the potential trade-offs when validation-test discrepancies exist.

Table 11: Performance comparison across different types of weight decay.

Dataset	Method	Utility (Acc or AUC)	Worst	Gap	EqOdd	DP
	No	86.52 ± 0.2	83.65 ± 0.3	6.91 ± 0.4	82.15 ± 1.4	67.69 ± 1.1
CelebA	L1	86.52 ± 0.2	83.68 ± 0.3	6.82 ± 0.4	81.48 ± 1.1	67.10 ± 1.0
	L2	86.57 ± 0.2	83.76 ± 0.2	6.76 ± 0.3	81.91 ± 0.6	67.20 ± 0.7
	No	93.11 ± 0.5	91.86 ± 0.9	2.92 ± 1.0	97.22 ± 1.0	94.69 ± 1.3
UTKFace	L1	93.02 ± 0.3	91.85 ± 0.5	2.71 ± 1.2	97.21 ± 1.1	95.13 ± 1.2
	L2	92.75 ± 0.5	91.78 ± 0.6	2.26 ± 0.6	97.62 ± 0.5	94.55 ± 1.2
	No	67.01 ± 0.2	66.71 ± 0.3	0.63 ± 0.5	95.99 ± 0.5	97.47 ± 0.4
FairFace	L1	66.60 ± 0.2	66.07 ± 0.3	1.13 ± 0.6	96.20 ± 0.4	97.61 ± 0.3
	L2	66.76 ± 0.2	66.34 ± 0.4	0.87 ± 0.5	96.22 ± 0.6	97.61 ± 0.4
	No	67.60 ± 0.4	64.24 ± 0.6	4.39 ± 0.9	96.26 ± 1.1	95.39 ± 1.0
Facet	L1	67.85 ± 0.2	63.71 ± 0.7	5.41 ± 0.9	96.84 ± 0.8	96.12 ± 0.7
	L2	67.55 ± 0.4	64.25 ± 1.0	4.31 ± 1.1	96.47 ± 1.2	95.40 ± 1.1
	No	88.12 ± 1.9	84.85 ± 2.0	3.42 ± 1.5	88.14 ± 4.1	83.78 ± 3.9
HAM10000	L1	87.66 ± 1.1	83.70 ± 3.1	3.86 ± 3.5	87.75 ± 4.0	83.95 ± 5.1
	L2	88.35 ± 1.8	84.68 ± 2.0	4.11 ± 2.1	88.17 ± 3.1	82.22 ± 4.8
	No	89.65 ± 1.0	88.30 ± 1.0	2.94 ± 1.1	95.08 ± 2.8	94.36 ± 1.3
Fitz17k	L1	82.26 ± 1.7	79.79 ± 2.5	4.70 ± 2.4	97.10 ± 1.4	97.91 ± 0.4
	L2	89.74 ± 1.0	88.39 ± 1.1	2.92 ± 1.1	94.92 ± 2.5	94.46 ± 1.4
	No	85.38 ± 0.7	83.38 ± 1.6	4.00 ± 2.0	69.10 ± 1.0	76.53 ± 1.6
Waterbirds	L1	77.46 ± 0.6	76.79 ± 1.1	1.33 ± 0.9	88.83 ± 2.3	91.68 ± 2.0
	L2	85.63 ± 1.4	84.20 ± 0.9	2.87 ± 0.9	66.53 ± 3.3	77.67 ± 3.5

Table 12: Performance comparison across different model sizes.

Dataset	Batch Size	Utility (Acc or AUC)	Worst	Gap	EqOdd	DP
	ResNet-18	92.75 ± 0.5	91.78 ± 0.6	2.26 ± 0.6	97.62 ± 0.5	94.55 ± 1.2
	ResNet-34	93.27 ± 0.8	92.31 ± 0.8	2.25 ± 0.4	97.66 ± 0.3	94.66 ± 1.0
UTKFace	ResNet-50	93.39 ± 0.7	92.50 ± 0.7	2.05 ± 0.8	97.85 ± 0.7	95.14 ± 1.9
UTKFace	ResNet-101	92.74 ± 0.0	91.11 ± 0.0	3.82 ± 0.0	96.23 ± 0.0	96.25 ± 0.0
	ResNet-18	66.76 ± 0.2	66.34 ± 0.4	0.87 ± 0.5	96.22 ± 0.6	97.61 ± 0.4
	ResNet-34	68.22 ± 0.2	67.80 ± 0.2	0.90 ± 0.5	96.33 ± 0.6	97.46 ± 0.3
FairFace	ResNet-50	67.33 ± 0.7	67.06 ± 0.6	0.58 ± 0.3	96.34 ± 0.6	97.57 ± 0.3
ranrace	ResNet-101	68.68 ± 0.6	68.31 ± 0.6	0.79 ± 0.3	95.97 ± 0.7	97.38 ± 0.4
	ResNet-18	88.35 ± 1.8	84.68 ± 2.0	4.11 ± 2.1	88.17 ± 3.1	82.22 ± 4.8
	ResNet-34	87.44 ± 0.3	83.35 ± 1.9	4.87 ± 2.1	88.33 ± 3.0	82.08 ± 4.7
HAM10000	ResNet-50	88.35 ± 0.9	83.94 ± 2.6	5.49 ± 2.8	84.33 ± 6.8	79.63 ± 6.8
HAMHUUUU	ResNet-101	89.01 ± 1.7	85.44 ± 3.5	4.69 ± 3.0	89.67 ± 3.3	84.84 ± 2.5

Table 13: Comparison of model selections. M denotes selecting models based on the maximal overall utility.

Dataset	CelebA					τ	TKFA	CE				FairFa	ce				Facet			
Dutuset	ACC	Worst	Gap	EqOdd	DP	ACC	Worst	Gap	EqOdd	DP	ACC	Worst	Gap	EqOdd	DP	ACC	Worst	Gap	EqOdd	DP
ERM-M	86.56	83.69	6.90	81.89	67.48	92.78	91.68	2.58	97.42	94.44	66.76	66.34	0.87	96.22	97.61	67.90	64.16	4.89	96.90	95.95
ERM-DTO	86.57	83.76	6.76	81.91	67.20	92.75	91.78	2.26	97.62	94.55	66.76	66.34	0.87	96.22	97.61	67.55	64.25	4.31	96.47	95.40
RandAug-M	86.7 3	83.89	6.82	81.16	66.90	93.31	92.17	2.64	97.32	95.11	68.50	67.70	1.70	96.17	97.57	67.97	64.55	4.46	96.81	96.02
RandAug-FWH	86.72	83.89	6.80	81.39	66.99	93.19	92.19	2.34	97.62	94.83	68.37	67.69	1.44	96.14	97.55	67.83	64.94	3.78	96.68	95.91
GroupDRO-M	86.14	83.51	6.32	78.76	66.32	92.37	91.43	2.18	97.46	94.46	65.81	65.03	1.64	96.53	97.66	67.32	64.19	4.09	93.57	93.16
GroupDRO-FWH	86.12	83.50	6.31	78.73	66.41	92.45	91.41	2.44	97.06	94.78	65.51	65.22	0.60	96.31	97.47	67.20	64.07	4.08	93.76	93.16
Dataset	HAM10000				Fitz17k				Waterbirds											
Dutuset	ACC	Worst	Gap	EqOdd	DP	ACC	Worst	Gap	EqOdd	DP	ACC	Worst	Gap	EqOdd	DP					
ERM-M	88.27	83.61	5.30	86.99	82.41	89.82	88.46	3.01	94.72	94.33	85.63	84.20	2.87	66.53	77.67					
ERM-DTO	88.35	84.67	4.11	86.79	81.48	89.74	88.39	2.92	94.92	94.46	85.63	84.20	2.87	66.53	77.67					
RandAug-M	89.53	85.33	4.63	86.34	79.94	91.24	90.13	2.52	95.36	94.48	87.48	85.60	3.76	72.12	80.35					
RandAug-FWH	89.09	84.67	4.99	88.43	84.73	91.29	90.15	2.51	95.61	94.53	86.09	84.52	3.14	68.99	77.25					
GroupDRO-M	88.96	85.23	4.33	86.12	79.45	91.04	89.66	3.55	95.22	94.43	86.78	84.76	4.05	71.88	80.88					
GroupDRO-FWH	87.66	83.98	4.98	90.94	83.86	90.72	90.06	1.92	94.36	95.24	85.46	84.45	2.02	67.31	76.91					

C.2 COMPUTATIONAL OVERHEAD COMPARISON

In this section, we compare the computational overhead of various methods in terms of additional model parameters and overall computation times. Since methods like **FIS** perform additional computations during validation, we measure the total computational time, including the training, validation, and testing phases. Notably, because **FIS** requires computing the gradient for each

sample, we use PyTorch's vmap to accelerate the computation, which aligns with the official **FIS** implementation in JAX.

In Table 14, we report the additional parameters required by specific methods and the average computational times (in seconds) for the evaluated methods. Most methods use the same number of parameters as the standard ERM baseline, so we focus only on those that introduce extra model parameters. The reported times are averaged over 5 runs, including the training, validation, and testing phases. For datasets with a limited number of samples, time differences are negligible and thus omitted from this comparison. Although some methods, like **LAFTR**, exhibit significantly increased computational demands, they do not outperform simpler methods. This highlights the need for easy yet effective approaches that balance performance with computational efficiency.

Table 14: Average computation time in seconds for one round and the extra number of parameters introduced by specific methods. DFR is not included here since it is a post-processing method.

	ERM	RandAug	Mixup	Resampling	BM	FIS
Parameters	_	_	_	_	1024	_
CelebA	113.06 ± 0.2	117.06 ± 1.9	112.48 ± 1.2	112.80 ± 0.4	114.08 ± 1.3	145.29 ± 0.5
UTKFace	39.33 ± 1.9	39.55 ± 1.0	38.44 ± 0.5	38.83 ± 1.1	39.79 ± 1.3	41.65 ± 0.6
FairFace	54.54 ± 0.3	54.53 ± 0.4	54.90 ± 0.3	54.68 ± 0.4	55.40 ± 0.4	72.07 ± 0.1
Facet	26.44 ± 0.1	26.51 ± 0.2	26.69 ± 0.2	26.53 ± 0.2	27.21 ± 0.2	35.14 ± 0.2
Average	58.34	59.42	58.13	58.21	59.12	73.54
	Decoupled	LAFTR	FSCL	GapReg	MCDP	GroupDRO
Parameters	1024	65921	_	_	_	_
CelebA	114.24 ± 0.6	175.77 ± 0.9	119.75 ± 2.1	112.92 ± 0.7	112.94 ± 1.0	113.75 ± 1.4
UTKFace	38.69 ± 1.0	40.56 ± 1.1	47.62 ± 0.4	38.94 ± 0.4	38.97 ± 0.8	39.00 ± 1.0
FairFace	55.38 ± 0.6	83.76 ± 0.4	54.81 ± 0.9	55.63 ± 0.5	55.63 ± 0.4	55.31 ± 0.7
Facet	26.74 ± 0.2	39.64 ± 0.2	26.19 ± 0.2	26.85 ± 0.3	26.82 ± 0.3	26.71 ± 0.2
Average	58.76	84.93	62.09	58.59	58.59	58.69

C.3 LARGE VISION-LANGUAGE MODEL EXPERIMENTS

C.3.1 MODEL SCALE IMPACT ON FAIRNESS

Table 15: Effect of LVLM model scale on accuracy and fairness metrics

Dataset	Metric	7B	LLaVA-1.6 13B	34B	7B	Qwen2.5-VL 32B	72B	4B	Gemma 3 12B	27B	3.2-11B	Llama 3.2-90B	4-Scout-109B
CelebA	ACC Worst Gap EqOdd DP	$\begin{array}{c} 50.21 \pm 0.2 \\ 46.79 \pm 0.6 \\ 5.85 \pm 0.7 \\ 88.81 \pm 0.3 \\ 77.63 \pm 0.2 \end{array}$	$\begin{array}{c} 46.50 \pm 0.1 \\ 38.93 \pm 0.3 \\ 12.95 \pm 0.4 \\ 92.44 \pm 0.2 \\ 83.16 \pm 0.2 \end{array}$	$\begin{array}{c} 44.83 \pm 0.0 \\ 32.69 \pm 0.3 \\ 20.75 \pm 0.5 \\ 97.41 \pm 0.2 \\ 91.92 \pm 0.2 \end{array}$	$\begin{array}{c} 65.02 \pm 0.2 \\ 63.41 \pm 0.3 \\ 2.75 \pm 0.5 \\ 95.27 \pm 0.8 \\ 81.51 \pm 0.9 \end{array}$	$\begin{array}{c} 42.58 \pm 0.2 \\ 27.79 \pm 0.2 \\ 25.33 \pm 0.6 \\ 95.26 \pm 0.5 \\ 99.26 \pm 0.1 \end{array}$	$68.76 \pm 0.1 \\ 66.02 \pm 0.4 \\ 4.70 \pm 0.7 \\ 92.69 \pm 0.3 \\ 91.71 \pm 0.2$	$\begin{array}{c} 45.08 \pm 0.2 \\ 37.15 \pm 0.2 \\ 13.57 \pm 0.4 \\ 92.25 \pm 0.2 \\ 83.41 \pm 0.4 \end{array}$	66.23 ± 0.2 65.17 ± 0.4 1.83 ± 0.4 89.72 ± 0.5 73.20 ± 0.4	$74.04 \pm 0.1 \\72.18 \pm 0.0 \\4.47 \pm 0.0 \\92.76 \pm 0.8 \\74.53 \pm 0.8$	$\begin{array}{c} 82.04 \pm 0.2 \\ 78.00 \pm 0.4 \\ 9.64 \pm 0.6 \\ 86.63 \pm 0.6 \\ 76.33 \pm 0.3 \end{array}$	$76.23 \pm 0.1 \\ 68.79 \pm 0.5 \\ 17.89 \pm 0.9 \\ 94.56 \pm 3.6 \\ 88.27 \pm 1.5$	$\begin{array}{c} 83.71 \pm 0.2 \\ 80.54 \pm 0.2 \\ 7.62 \pm 0.4 \\ 84.81 \pm 0.9 \\ 71.52 \pm 0.5 \end{array}$
UTKFace	ACC Worst Gap EqOdd DP	$\begin{array}{c} 96.76 \pm 0.3 \\ 95.71 \pm 0.3 \\ 2.44 \pm 0.6 \\ 97.51 \pm 0.6 \\ 95.15 \pm 1.3 \end{array}$	$\begin{array}{c} 97.01 \pm 0.0 \\ 96.28 \pm 0.1 \\ 1.71 \pm 0.4 \\ 98.31 \pm 0.4 \\ 94.97 \pm 1.6 \end{array}$	$\begin{array}{c} 97.12 \pm 0.2 \\ 96.34 \pm 0.1 \\ 1.81 \pm 0.2 \\ 98.16 \pm 0.2 \\ 95.22 \pm 1.3 \end{array}$	$\begin{array}{c} 96.12 \pm 0.1 \\ 95.03 \pm 0.2 \\ 2.51 \pm 0.7 \\ 97.09 \pm 0.6 \\ 93.05 \pm 1.4 \end{array}$	$\begin{array}{c} 97.61 \pm 0.2 \\ 96.86 \pm 0.1 \\ 1.74 \pm 0.2 \\ 98.19 \pm 0.2 \\ 93.96 \pm 1.6 \end{array}$	$\begin{array}{c} 96.78 \pm 0.4 \\ 95.76 \pm 0.4 \\ 2.36 \pm 0.8 \\ 97.59 \pm 0.8 \\ 94.33 \pm 1.6 \end{array}$	$\begin{array}{c} 96.33 \pm 0.3 \\ 95.41 \pm 0.2 \\ 2.15 \pm 0.5 \\ 97.68 \pm 0.5 \\ 93.54 \pm 1.4 \end{array}$	$\begin{array}{c} 97.25\pm0.2\\ 96.70\pm0.1\\ 1.29\pm0.4\\ 98.71\pm0.4\\ 94.84\pm1.1 \end{array}$	$\begin{array}{c} 97.25 \pm 0.5 \\ 96.89 \pm 0.4 \\ 0.85 \pm 0.2 \\ 99.20 \pm 0.2 \\ 95.28 \pm 1.1 \end{array}$	$\begin{array}{c} 96.81 \pm 0.3 \\ 96.04 \pm 0.2 \\ 1.81 \pm 0.3 \\ 98.03 \pm 0.3 \\ 93.53 \pm 1.1 \end{array}$	$\begin{array}{c} 97.28 \pm 0.1 \\ 96.31 \pm 0.0 \\ 2.26 \pm 0.4 \\ 97.73 \pm 0.4 \\ 94.18 \pm 1.3 \end{array}$	97.02 ± 0.2 96.29 ± 0.2 1.70 ± 0.3 98.27 ± 0.3 94.90 ± 1.3
FairFace	ACC Worst Gap EqOdd DP	$\begin{array}{c} 53.07 \pm 0.0 \\ 49.98 \pm 0.0 \\ 6.55 \pm 0.0 \\ 95.04 \pm 0.0 \\ 96.34 \pm 0.0 \end{array}$	$\begin{array}{c} 59.59 \pm 0.0 \\ 58.54 \pm 0.0 \\ 1.97 \pm 0.0 \\ 94.55 \pm 0.0 \\ 95.64 \pm 0.0 \end{array}$	$66.91 \pm 0.0 \\ 66.04 \pm 0.0 \\ 1.84 \pm 0.0 \\ 96.22 \pm 0.0 \\ 98.07 \pm 0.0$	$\begin{array}{c} 66.09 \pm 0.0 \\ 65.07 \pm 0.0 \\ 2.17 \pm 0.0 \\ 96.04 \pm 0.0 \\ 97.21 \pm 0.0 \end{array}$	$63.98 \pm 0.0 \\ 62.65 \pm 0.0 \\ 2.83 \pm 0.0 \\ 95.50 \pm 0.0 \\ 97.23 \pm 0.0$	$65.57 \pm 0.0 \\ 64.73 \pm 0.0 \\ 1.79 \pm 0.0 \\ 96.26 \pm 0.0 \\ 97.57 \pm 0.0$	$\begin{array}{c} 52.73 \pm 0.1 \\ 50.55 \pm 0.1 \\ 4.63 \pm 0.1 \\ 96.01 \pm 0.1 \\ 97.50 \pm 0.0 \end{array}$	$\begin{array}{c} 57.27 \pm 2.8 \\ 55.35 \pm 3.0 \\ 4.08 \pm 0.3 \\ 93.07 \pm 1.9 \\ 95.80 \pm 1.1 \end{array}$	$\begin{array}{c} 59.40 \pm 0.1 \\ 56.71 \pm 0.0 \\ 5.71 \pm 0.1 \\ 95.13 \pm 0.1 \\ 97.16 \pm 0.1 \end{array}$	$\begin{array}{c} 51.70 \pm 0.1 \\ 49.73 \pm 0.3 \\ 4.18 \pm 0.3 \\ 95.99 \pm 0.1 \\ 96.76 \pm 0.0 \end{array}$	$\begin{array}{c} 61.26 \pm 0.2 \\ 60.96 \pm 0.3 \\ 0.63 \pm 0.2 \\ 95.55 \pm 0.2 \\ 96.58 \pm 0.0 \end{array}$	56.52 ± 0.2 53.53 ± 0.0 6.35 ± 0.4 96.54 ± 0.1 98.00 ± 0.1
Facet	ACC Worst Gap EqOdd DP	$\begin{array}{c} 63.50 \pm 0.2 \\ 62.41 \pm 0.3 \\ 1.42 \pm 0.2 \\ 94.98 \pm 1.6 \\ 93.28 \pm 1.4 \end{array}$	66.80 ± 0.3 62.90 ± 1.1 5.09 ± 1.3 98.19 ± 0.9 96.79 ± 1.1	$\begin{array}{c} 58.14 \pm 0.2 \\ 57.97 \pm 0.4 \\ 0.72 \pm 0.6 \\ 95.90 \pm 1.4 \\ 93.91 \pm 1.5 \end{array}$	$\begin{array}{c} 67.91 \pm 0.5 \\ 64.06 \pm 1.2 \\ 5.02 \pm 1.0 \\ 96.29 \pm 0.9 \\ 95.81 \pm 0.9 \end{array}$	67.80 ± 0.3 63.16 ± 1.3 6.06 ± 1.6 97.67 ± 0.7 97.22 ± 0.6	$68.41 \pm 0.5 \\ 63.79 \pm 1.3 \\ 6.04 \pm 1.2 \\ 97.04 \pm 1.2 \\ 96.15 \pm 1.4$	$\begin{array}{c} 66.93 \pm 0.4 \\ 62.39 \pm 0.9 \\ 5.92 \pm 0.9 \\ 97.88 \pm 0.4 \\ 98.28 \pm 0.9 \end{array}$	$\begin{array}{c} 54.47 \pm 0.2 \\ 53.82 \pm 0.5 \\ 2.80 \pm 1.8 \\ 97.63 \pm 0.4 \\ 96.13 \pm 0.7 \end{array}$	$\begin{array}{c} 57.75 \pm 0.7 \\ 57.25 \pm 0.5 \\ 1.42 \pm 1.2 \\ 96.22 \pm 1.0 \\ 94.56 \pm 0.6 \end{array}$	$66.74 \pm 0.4 \\ 61.23 \pm 0.5 \\ 7.19 \pm 0.9 \\ 93.68 \pm 1.8 \\ 95.72 \pm 1.7$	$67.75 \pm 0.5 \\ 63.05 \pm 1.1 \\ 6.12 \pm 0.8 \\ 97.23 \pm 1.1 \\ 96.05 \pm 1.1$	67.54 ± 0.4 64.61 ± 1.0 3.83 ± 1.0 97.55 ± 1.1 96.28 ± 1.0
Waterbirds	ACC Worst Gap EqOdd DP	$\begin{array}{c} 77.24 \pm 0.0 \\ 75.63 \pm 0.0 \\ 3.21 \pm 0.0 \\ 63.22 \pm 0.0 \\ 72.83 \pm 0.0 \end{array}$	79.46 ± 0.0 76.08 ± 0.0 6.77 ± 0.0 65.54 ± 0.0 72.45 ± 0.0	$\begin{array}{c} 86.95 \pm 0.0 \\ 81.26 \pm 0.0 \\ 11.39 \pm 0.0 \\ 80.88 \pm 0.0 \\ 80.46 \pm 0.0 \end{array}$	$\begin{array}{c} 91.68 \pm 0.0 \\ 90.68 \pm 0.0 \\ 2.00 \pm 0.0 \\ 90.21 \pm 0.0 \\ 92.13 \pm 0.0 \end{array}$	$\begin{array}{c} 91.62 \pm 0.0 \\ 90.46 \pm 0.0 \\ 2.32 \pm 0.0 \\ 91.90 \pm 0.0 \\ 93.12 \pm 0.0 \end{array}$	$\begin{array}{c} 92.41 \pm 0.0 \\ 91.27 \pm 0.0 \\ 2.28 \pm 0.0 \\ 94.03 \pm 0.0 \\ 94.62 \pm 0.0 \end{array}$	$\begin{array}{c} 75.59 \pm 1.8 \\ 62.40 \pm 4.2 \\ 26.39 \pm 4.9 \\ 61.61 \pm 0.7 \\ 58.82 \pm 3.2 \end{array}$	$\begin{array}{c} 92.20\pm0.0 \\ 91.98\pm0.0 \\ 0.44\pm0.1 \\ 89.28\pm0.2 \\ 92.36\pm0.1 \end{array}$	$\begin{array}{c} 90.77 \pm 0.0 \\ 88.97 \pm 0.2 \\ 3.59 \pm 0.3 \\ 92.99 \pm 0.5 \\ 93.17 \pm 0.1 \end{array}$	$\begin{array}{c} 75.41 \pm 0.3 \\ 61.61 \pm 0.6 \\ 27.59 \pm 0.5 \\ 65.72 \pm 0.6 \\ 60.97 \pm 0.6 \end{array}$	$77.62 \pm 0.2 \\ 63.35 \pm 0.3 \\ 28.55 \pm 0.3 \\ 71.54 \pm 0.6 \\ 64.46 \pm 0.4$	$\begin{array}{c} 88.62 \pm 0.1 \\ 85.74 \pm 0.2 \\ 5.76 \pm 0.2 \\ 89.38 \pm 0.1 \\ 89.44 \pm 0.1 \end{array}$
Average	ACC Worst Gap EqOdd DP	68.156 66.104 3.894 87.912 87.046	69.872 66.546 5.698 89.806 88.602	70.79 66.86 7.302 93.714 91.916	77.364 75.65 2.89 94.98 91.942	72.718 68.184 7.656 95.704 96.158	78.386 76.314 3.434 95.522 94.876	67.332 61.58 10.532 89.086 86.31	73.484 72.604 2.088 93.682 90.466	75.842 74.4 3.208 95.26 90.94	74.54 69.322 10.082 88.01 84.662	76.006 70.51 11.002 91.416 87.934	78.682 76.142 5.052 93.31 90.028

In this paper, we include four families of vision—language foundation models: LLaVA, Qwen-VL, Gemma, and Llama at multiple parameter scales (from roughly 4 billion to 90 billion weights) to probe whether the standard "bigger-is-better" rule extends to fairness (Table 15). The averaged results

Table 16: Prompts used for evaluating LVLM sensitivity. Each dataset has the original instruction (Origin) and two paraphrased variants.

Dataset	Prompts
CelebA	Origin: Is the person in the photo wavy-haired? Answer 'Yes' for wavy hair, 'No' for non-wavy hair. Variant 1: Does this person appear to have wavy hair? Reply with Yes/No only. Variant 2: Classify hair texture as wavy or non-wavy. Output: Yes or No.
FairFace	Origin: Based on the photo, classify the race of the person into: 0: White, 1: Black, 2: Latino or Hispanic, 3: East Asian, 4: Southeast Asian, 5: Indian, 6: Middle Eastern. Answer with the category number only. Variant 1: Pick a single integer label (0–6) for race: 0-White, 1-Black, 2-Latino/Hispanic, 3-East Asian, 4-Southeast Asian, 5-Indian, 6-Middle Eastern. Output digits only. Variant 2: Classify race using numbers: 0=White, 1=Black, 2=Latino/Hispanic, 3=East Asian, 4=Southeast Asian, 5=Indian, 6=Middle Eastern. Answer with number only.
Facet	Origin: Does the photo show a visible face? Answer 'Yes' if the face is visible, otherwise 'No'. Variant 1: Is a face clearly visible? Reply Yes/No only. Variant 2: Face visibility check: Yes or No? Output one word.

confirm that scale almost always boosts utility: mean accuracy climbs steadily within each family (e.g., LLaVA $68 \rightarrow 71$ %, Gemma $67 \rightarrow 76$ %), and worst-group accuracy rises in the same trend. Larger models also score better on our parity metrics like demographic parity (DP) and equalized odds (EqOdd), which suggests that they become more confident and consistent across groups. However, larger models do not guarantee a better accuracy gap. For example, the gap more than doubles ($3.9 \rightarrow 7.3$ %) for LLaVA. In short, increasing size is a reliable path to higher accuracy and better parity rates, but it does not guarantee a smaller disparity between groups; in some cases, it even amplifies it. These mixed trends highlight that parameter count alone cannot close fairness gaps where architectural choices, alignment strategies, and dataset biases remain decisive factors.

We further select the largest variant from each model family and present the results in Figure 8. Among them, Qwen2.5-VL 72B demonstrates the best overall balance between fairness and utility. Notably, Gemma 3 27B achieves comparable performance to Qwen despite its smaller model size, while LLaMA 4-Scout, although achieving strong headline accuracy, remains susceptible to dataset-specific biases.

C.3.2 PROMPT SENSITIVITY OF LVLMS

Fairness evaluations of LVLMs often rely on handcrafted prompts, raising concerns about robustness to prompt variations. To examine the sensitivity of LVLMs to prompt wording, we design two paraphrased variants for three dataset tasks (we consider those with relatively larger disparities) in addition to the original instruction (Table 16), and evaluate whether performance holds consistent across different formulations. Results in Table 17 show that, across most cases, models exhibit broadly consistent fairness performance under different prompts, with only a few extreme cases (e.g., Qwen2.5-VL on CelebA) showing large deviations. This indicates that prompt rewording alone does not fundamentally resolve fairness issues, LVLMs remain subject to subgroup disparities regardless of the prompt design.

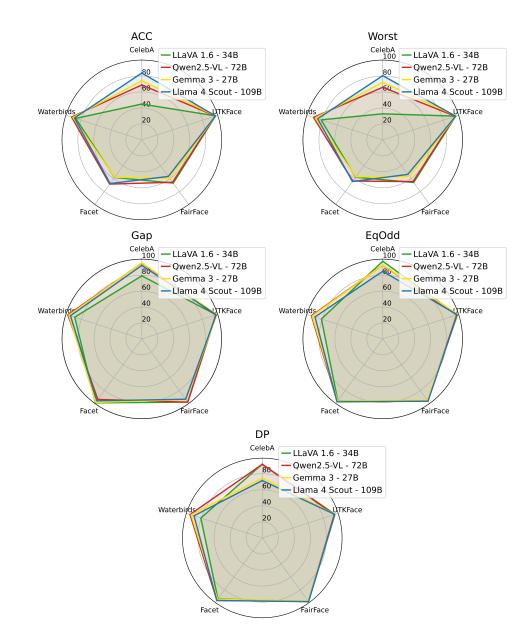


Figure 8: Performance comparison of LVLMs across five datasets and five evaluation metrics. For clarity, the Gap metric is shown as 100-Gap, where higher values indicate better performance.

Table 17: LVLM performance under different prompts. Each model is evaluated with the original instruction and two paraphrased variants. For CelebA, Llama3.2 often produces non-conforming answers under Variant 2 instead of the required categorical outputs; it is therefore omitted from the reported results.

Dataset	Metric	1	LLaVA-1.6-13I	3	l Q	wen2.5-VL-32	В		Gemma3-12B		l	Llama3.2-11B	
Dataset	Metric	Origin	Variant 1	Variant 2									
	ACC	46.50 ± 0.1	53.84 ± 0.2	50.30 ± 0.1	42.58 ± 0.2	33.47 ± 0.1	37.08 ± 0.0	66.23 ± 0.2	63.76 ± 0.2	74.77 ± 0.1	76.23 ± 0.1	81.34 ± 0.2	-
	Worst	38.93 ± 0.3	46.81 ± 0.4	45.10 ± 0.4	27.79 ± 0.2	15.16 ± 0.2	18.22 ± 0.0	65.17 ± 0.4	61.51 ± 0.4	64.98 ± 0.2	68.79 ± 0.5	77.10 ± 0.4	-
CelebA	Gap	12.95 ± 0.4	12.03 ± 0.4	8.89 ± 0.6	25.33 ± 0.6	31.32 ± 0.6	32.30 ± 0.0	1.83 ± 0.4	3.86 ± 0.3	23.58 ± 0.3	17.89 ± 0.9	10.20 ± 0.5	-
	EqOdd	92.44 ± 0.2	94.23 ± 0.2	91.04 ± 0.3	95.26 ± 0.5	98.61 ± 0.1	93.73 ± 0.0	89.72 ± 0.5	90.24 ± 0.3	93.32 ± 0.2	94.56 ± 3.6	92.56 ± 0.4	-
	DP	83.16 ± 0.2	87.06 ± 0.4	80.21 ± 0.2	99.26 ± 0.1	99.00 ± 0.1	94.81 ± 0.0	73.20 ± 0.4	74.62 ± 0.2	92.66 ± 0.6	88.27 ± 1.5	79.98 ± 0.6	-
	ACC	59.59 ± 0.0	61.57 ± 0.0	59.17 ± 0.0	63.98 ± 0.0	63.13 ± 0.0	63.00 ± 0.0	57.27 ± 2.8	57.15 ± 0.1	58.19 ± 0.1	61.26 ± 0.2	31.19 ± 0.1	40.66 ± 0.2
	Worst	58.54 ± 0.0	58.98 ± 0.0	58.01 ± 0.0	62.65 ± 0.0	61.41 ± 0.0	61.43 ± 0.0	55.35 ± 3.0	54.97 ± 0.0	56.09 ± 0.1	60.96 ± 0.3	28.99 ± 0.3	38.88 ± 0.5
FairFace	Gap	1.97 ± 0.0	5.49 ± 0.0	2.47 ± 0.0	2.83 ± 0.0	3.65 ± 0.0	3.33 ± 0.1	4.08 ± 0.3	4.64 ± 0.1	4.45 ± 0.1	0.63 ± 0.2	4.67 ± 0.6	3.78 ± 0.7
	EqOdd	94.55 ± 0.0	97.04 ± 0.0	96.82 ± 0.0	95.50 ± 0.0	97.11 ± 0.0	95.86 ± 0.0	93.07 ± 1.9	92.13 ± 0.0	92.12 ± 0.1	95.55 ± 0.2	97.11 ± 0.2	95.15 ± 0.1
	DP	95.64 ± 0.0	97.54 ± 0.0	97.06 ± 0.0	97.23 ± 0.0	97.75 ± 0.0	97.56 ± 0.0	95.80 ± 1.1	95.15 ± 0.1	95.38 ± 0.0	96.58 ± 0.0	98.69 ± 0.1	97.19 ± 0.1
	ACC	66.80 ± 0.3	66.63 ± 0.2	55.20 ± 0.5	67.80 ± 0.3	67.19 ± 0.1	67.59 ± 0.3	54.47 ± 0.2	55.51 ± 0.3	55.88 ± 0.5	67.75 ± 0.5	63.69 ± 0.1	38.03 ± 0.0
	Worst	62.90 ± 1.1	63.26 ± 0.5	54.28 ± 0.7	63.16 ± 1.3	62.40 ± 1.5	62.41 ± 1.3	53.82 ± 0.5	54.91 ± 0.1	55.08 ± 0.3	63.05 ± 1.1	60.03 ± 0.5	36.65 ± 0.2
Facet	Gap	5.09 ± 1.3	4.38 ± 0.8	3.95 ± 0.9	6.06 ± 1.6	6.24 ± 1.8	6.75 ± 1.3	2.80 ± 1.8	2.57 ± 1.1	1.41 ± 0.5	6.12 ± 0.8	4.75 ± 0.7	6.02 ± 0.8
	EqOdd	98.19 ± 0.9	98.16 ± 0.5	97.19 ± 0.8	97.67 ± 0.7	97.27 ± 0.4	96.70 ± 0.0	97.63 ± 0.4	96.97 ± 1.3	95.77 ± 1.6	97.23 ± 1.1	96.24 ± 1.1	98.79 ± 0.2
	DP	96.79 ± 1.1	96.79 ± 1.4	98.92 ± 0.3	97.22 ± 0.6	96.87 ± 0.3	95.67 ± 0.2	96.13 ± 0.7	95.60 ± 1.3	93.87 ± 1.8	96.05 ± 1.1	97.24 ± 0.9	99.48 ± 0.4

D ADDING NEW DATASETS AND ALGORITHMS TO THE FRAMEWORK

Our benchmarking framework provides a flexible and extensible design for integrating new image classification datasets. A new dataset can be inherited from FairDataset and specify image folder location, target labels, sensitive attributes, and an indices list to split the dataset. The core dataset class, FairDataset, is presented as below.

```
class FairDataset(Dataset):
    def __init__(self, root, split='train', transform=None, seed=42):
        # Base folder and images directory
        self.root = root
        self.img_dir = None

        self.split = split
        self.indices = [] # Indices for the split dataset
        self.sensitive_attrs = np.array([]) # Sensitive attributes
        self.targets = np.array([]) # Target labels

        self.transform = transform # Image transformation pipeline
        self.image_file_list = []

# Option to store images in memory for faster access
        self.images = []
```

By inheriting from ERM, this new class has access to the default model, optimizer, and loss function. A new method needs to be implemented to rewrite the training process, ensuring that it aligns with the new algorithm's objectives. For example, the MCDP class extends ERM and introduces a fairness-aware loss function.

```
1810
      class MCDP (erm):
         def __init__(self, args):
1811
             super().__init__(args)
1812
             self.fair_loss = MaxCDFdp(args.mcdp_temperature)
1813
         def train(self, train_loader, epoch, args):
             self.model.train()
1815
             for batch_index, (data, target, sensitive_attr) in
1816
              1817
                 output = self.model(data)
1818
                 self.optimizer.zero_grad()
1819
                 fairloss = args.mcdp_lambda * self.fair_loss(output,
1820
                 loss = self.criterion(output, target) + fairloss
1821
                 loss.backward()
1822
                 self.optimizer.step()
```

E SUPPLEMENTARY RESULTS

Here we present the standard deviation for Table 2 and Table 19 below:

Table 18: Full results on seven datasets with standard deviations for vision models.

Dataset	Metric	ERM	RandAug	Mixup	Resampling	ВМ	FIS	Decoupled	LAFTR	FSCL	GapReg	MCDP	GroupDRO	DFR
	ACC	86.57 ± 0.2	86.72 ± 0.2	85.61 ± 0.3	86.35 ± 0.2	85.93 ± 0.1	83.05 ± 0.2	86.35 ± 0.2	86.55 ± 0.2	85.61 ± 0.1	85.62 ± 0.6	80.26 ± 0.8	86.12 ± 0.2	86.58 ± 0.1
	Worst	83.76 ± 0.2	83.89 ± 0.3	82.74 ± 0.4	83.44 ± 0.2	82.86 ± 0.3	79.33 ± 0.5	83.46 ± 0.2	83.67 ± 0.3	82.56 ± 0.3	83.17 ± 0.4	77.13 ± 0.8	83.50 ± 0.2	83.78 ± 0.2
CelebA	Gap	6.76 ± 0.3	6.80 ± 0.4	6.90 ± 0.3	6.98 ± 0.3	7.38 ± 0.5	8.94 ± 0.7	6.93 ± 0.4	6.93 ± 0.5	7.35 ± 0.6	5.90 ± 0.9	7.52 ± 1.2	6.31 ± 0.5	6.74 ± 0.4
CCICDA	EqOdd	81.91 ± 0.6	81.73 ± 1.1	87.08 ± 2.2	81.80 ± 0.7	78.84 ± 1.1	75.79 ± 1.8	80.59 ± 1.3	81.15 ± 1.7	85.45 ± 1.4	93.94 ± 5.8	89.63 ± 3.3	78.73 ± 1.3	81.83 ± 0.7
	DP	67.20 ± 0.7	67.37 ± 0.4	70.83 ± 1.8	67.39 ± 0.3	66.91 ± 1.0	66.84 ± 1.8	66.68 ± 0.9	66.90 ± 1.2	69.72 ± 1.6	75.91 ± 4.5	93.11 ± 2.2	66.41 ± 1.2	67.30 ± 0.7
	ACC	92.75 ± 0.5	93.19 ± 0.3	92.62 ± 0.6	92.70 ± 0.5	93.33 ± 0.6	91.97 ± 0.5	91.68 ± 1.1	93.17 ± 0.3	93.52 ± 0.5	92.53 ± 0.9	92.49 ± 0.7	92.45 ± 0.4	92.73 ± 0.6
	Worst	91.78 ± 0.6	92.19 ± 0.3	91.55 ± 0.7	91.60 ± 0.7	92.27 ± 0.8	90.91 ± 0.5	90.84 ± 0.9	92.05 ± 0.3	92.62 ± 0.4	91.70 ± 0.9	91.63 ± 1.1	91.41 ± 0.6	91.60 ± 0.9
UTKFace	Gap	2.26 ± 0.6	2.34 ± 0.6	2.50 ± 0.4	2.56 ± 0.4	2.47 ± 0.5	2.48 ± 0.4	1.97 ± 1.0	2.61 ± 0.3	2.10 ± 0.8	1.91 ± 0.6	2.00 ± 1.0	2.44 ± 1.0	2.63 ± 0.8
UIKrace	EqOdd	97.62 ± 0.5	97.62 ± 0.4	97.61 ± 0.4	97.49 ± 0.4	97.39 ± 0.5	97.51 ± 0.4	97.43 ± 0.4	97.44 ± 0.3	97.44 ± 1.3	98.10 ± 0.6	98.04 ± 1.0	97.06 ± 0.6	97.39 ± 0.9
	DP	94.55 ± 1.2	94.83 ± 1.1	94.51 ± 1.5	95.34 ± 1.8	94.34 ± 1.5	94.69 ± 1.7	96.27 ± 2.4	95.44 ± 2.1	94.18 ± 2.1	95.30 ± 1.6	95.80 ± 0.6	94.78 ± 1.8	94.83 ± 1.4
	ACC	66.76 ± 0.2	68.37 ± 0.3	65.40 ± 0.6	65.40 ± 0.2	65.66 ± 1.1	65.31 ± 0.5	67.03 ± 0.5	66.44 ± 0.8	65.42 ± 0.5	65.02 ± 1.0	66.06 ± 1.1	65.51 ± 0.5	63.20 ± 2.1
	Worst	66.34 ± 0.4	67.69 ± 0.3	64.50 ± 0.6	64.51 ± 0.4	65.20 ± 1.0	64.59 ± 0.4	66.61 ± 0.5	65.60 ± 0.9	64.64 ± 0.4	64.12 ± 1.3	65.62 ± 0.9	65.22 ± 0.5	62.45 ± 2.5
FairFace	Gap	0.87 ± 0.5	1.44 ± 0.6	1.93 ± 0.3	1.90 ± 0.6	0.97 ± 0.2	1.53 ± 0.6	0.87 ± 0.5	1.76 ± 0.4	1.66 ± 0.5	1.92 ± 0.9	0.90 ± 0.5	0.60 ± 0.6	1.59 ± 0.7
ranrace	EqOdd	96.22 ± 0.6	96.14 ± 0.1	96.55 ± 0.3	96.83 ± 0.3	95.70 ± 0.3	95.88 ± 0.3	95.14 ± 0.7	94.73 ± 1.0	97.05 ± 0.5	96.15 ± 0.1	97.85 ± 0.2	96.31 ± 0.7	95.81 ± 0.3
	DP	97.61 ± 0.4	97.55 ± 0.2	97.62 ± 0.2	97.80 ± 0.2	97.30 ± 0.2	97.40 ± 0.2	96.88 ± 0.4	96.86 ± 0.6	98.10 ± 0.1	97.51 ± 0.2	98.50 ± 0.2	97.47 ± 0.4	97.43 ± 0.1
	ACC	67.55 ± 0.4	67.83 ± 0.7	67.86 ± 0.4	67.56 ± 0.2	65.87 ± 1.2	67.60 ± 0.4	67.33 ± 0.7	70.74 ± 7.5	67.79 ± 0.4	67.01 ± 0.2	67.91 ± 0.4	67.20 ± 0.2	66.87 ± 0.5
	Worst	64.25 ± 1.0	64.94 ± 1.0	64.54 ± 1.0	64.13 ± 0.5	62.67 ± 0.9	63.53 ± 0.9	62.60 ± 1.1	68.60 ± 8.6	65.02 ± 0.9	62.22 ± 1.4	64.21 ± 1.2	64.07 ± 0.8	63.10 ± 1.5
Facet	Gap	4.31 ± 1.1	3.78 ± 1.0	4.33 ± 1.4	4.48 ± 0.4	4.18 ± 0.6	5.33 ± 1.3	6.17 ± 0.6	2.82 ± 1.4	3.61 ± 0.8	6.26 ± 1.7	4.84 ± 1.5	4.08 ± 0.9	4.92 ± 1.9
rucci	EqOdd	96.47 ± 1.2	96.68 ± 1.0	97.50 ± 1.2	94.37 ± 0.8	96.32 ± 1.8	98.10 ± 1.3	88.04 ± 1.9	96.38 ± 2.8	96.50 ± 1.6	98.92 ± 1.0	98.23 ± 1.0	93.76 ± 1.0	96.82 ± 1.0
	DP	95.40 ± 1.1	95.91 ± 1.0	96.71 ± 1.2	93.96 ± 0.7	95.09 ± 1.9	97.84 ± 1.2	87.30 ± 1.7	95.00 ± 3.0	95.66 ± 1.4	98.73 ± 1.4	98.46 ± 1.3	93.16 ± 1.2	96.09 ± 1.1
	AUC	88.35 ± 1.8	89.09 ± 1.3	86.51 ± 2.4	87.75 ± 2.2	89.54 ± 1.4	85.97 ± 1.1	87.87 ± 2.0	85.72 ± 2.7	86.12 ± 2.9	84.97 ± 3.2	82.96 ± 1.8	87.66 ± 2.7	87.06 ± 0.9
	Worst	84.68 ± 2.0	84.67 ± 3.3	82.31 ± 2.3	84.77 ± 3.3	86.49 ± 1.5	82.98 ± 1.6	84.04 ± 3.6	76.21 ± 4.2	78.36 ± 3.9	82.57 ± 3.7	80.29 ± 1.7	83.98 ± 3.0	82.49 ± 2.5
HAM10000	Gap	4.11 ± 2.1	4.99 ± 3.7	4.14 ± 2.8	3.52 ± 2.7	3.04 ± 0.9	3.11 ± 1.9	5.17 ± 2.7	13.21 ± 6.8	10.81 ± 3.9	3.07 ± 2.3	3.10 ± 2.2	4.98 ± 2.5	5.30 ± 2.2
11.1.1110000	EqOdd	88.17 ± 3.1	88.43 ± 4.0	91.14 ± 5.4	91.84 ± 3.3	85.65 ± 4.1	94.05 ± 2.3	77.52 ± 14.2	93.28 ± 6.2	84.57 ± 2.3	98.15 ± 3.7	99.52 ± 1.0	90.94 ± 5.3	91.25 ± 5.0
	DP	82.22 ± 4.8	84.73 ± 3.4	88.54 ± 8.9	85.05 ± 2.2	78.41 ± 4.9	88.58 ± 2.7	75.15 ± 11.6	91.00 ± 6.4	77.64 ± 4.5	96.74 ± 6.5	99.58 ± 0.8	83.86 ± 5.7	84.65 ± 4.6
	AUC	89.74 ± 1.0	91.29 ± 0.5	90.62 ± 0.9	90.76 ± 0.7	91.02 ± 0.6	88.34 ± 0.8	89.63 ± 0.5	91.50 ± 0.9	90.83 ± 1.9	89.59 ± 1.1	91.65 ± 0.7	90.72 ± 0.9	89.99 ± 0.9
	Worst	88.39 ± 1.1	90.15 ± 1.0	89.38 ± 1.9	88.99 ± 2.0	89.93 ± 0.8	87.02 ± 1.1	88.45 ± 1.3	90.59 ± 1.2	89.64 ± 2.2	88.52 ± 2.6	90.49 ± 0.8	90.06 ± 0.9	88.57 ± 0.5
Fitz17k	Gap	2.92 ± 1.1	2.51 ± 1.0	2.43 ± 1.4	3.62 ± 2.2	2.34 ± 1.4	3.06 ± 2.0	2.55 ± 2.2	2.95 ± 1.1	3.84 ± 1.2	1.84 ± 2.0	2.87 ± 1.1	1.92 ± 1.1	2.93 ± 1.5
FILZITA	EqOdd	94.92 ± 2.5	95.61 ± 2.3	96.20 ± 3.1	95.27 ± 3.3	94.99 ± 1.9	95.17 ± 2.4	94.09 ± 3.3	93.68 ± 3.4	97.45 ± 0.8	95.47 ± 1.5	95.68 ± 2.7	94.36 ± 2.4	94.30 ± 1.1
	DP	94.46 ± 1.4	94.53 ± 1.1	94.51 ± 1.4	93.31 ± 2.3	93.66 ± 2.1	95.60 ± 1.9	94.06 ± 1.9	94.46 ± 1.3	95.32 ± 1.4	96.42 ± 2.2	96.14 ± 2.1	95.24 ± 1.6	95.26 ± 1.2
	ACC	85.63 ± 1.4	86.09 ± 1.4	87.67 ± 1.1	87.35 ± 0.8	88.20 ± 1.0	83.72 ± 0.8	74.64 ± 3.0	85.72 ± 1.5	86.83 ± 2.1	86.45 ± 3.9	85.98 ± 0.9	85.46 ± 1.5	89.83 ± 1.4
	Worst	84.20 ± 0.9	84.52 ± 2.6	85.99 ± 1.5	84.85 ± 0.7	85.96 ± 1.1	82.67 ± 2.1	64.45 ± 5.7	83.94 ± 2.1	86.28 ± 2.3	85.72 ± 4.1	84.83 ± 1.4	84.45 ± 1.7	89.09 ± 1.7
Waterbirds	Gap	2.87 ± 0.9	3.14 ± 2.5	3.36 ± 1.7	4.98 ± 0.3	4.48 ± 0.6	2.09 ± 2.9	20.38 ± 5.6	3.56 ± 1.2	1.10 ± 0.6	1.47 ± 0.9	2.31 ± 1.4	2.02 ± 1.4	1.47 ± 0.6
···mcrbirds	EqOdd	66.53 ± 3.3	68.99 ± 2.5	81.42 ± 4.4	90.87 ± 1.3	77.21 ± 3.3	65.89 ± 2.1	47.31 ± 4.1	68.22 ± 3.0	90.00 ± 1.8	87.48 ± 8.8	72.97 ± 3.4	67.31 ± 3.8	97.79 ± 0.7
	DP	77.67 ± 3.5	77.25 ± 3.0	86.00 ± 4.6	90.93 ± 1.0	81.78 ± 2.3	75.30 ± 1.8	52.30 ± 5.8	77.47 ± 4.1	92.53 ± 1.4	91.39 ± 7.1	80.70 ± 1.8	76.91 ± 3.6	98.61 ± 0.9

Table 19: Full results on seven datasets with standard deviations for multimodal models..

Dataset	Metric	ERM	RandAug	BLIP2	CLIP	Fairer- CLIP	CLIP- SFID	LLaVA-1.6 34B	Qwen2.5-VL 72B	Gemma3 27B	Llama4 Scout
CelebA	ACC Worst Gap EqOdd DP	$86.57 \pm 0.2 \\ 83.76 \pm 0.2 \\ 6.76 \pm 0.3 \\ 81.91 \pm 0.6 \\ 67.20 \pm 0.7$	$86.72 \pm 0.2 83.89 \pm 0.3 6.80 \pm 0.4 81.73 \pm 1.1 67.37 \pm 0.4$	$\begin{array}{c} 47.38 \pm 0.2 \\ 36.82 \pm 0.6 \\ 18.09 \pm 0.7 \\ 97.24 \pm 0.6 \\ 95.90 \pm 0.7 \end{array}$	$74.07 \pm 0.4 \\ 67.43 \pm 0.5 \\ 15.97 \pm 0.4 \\ 83.72 \pm 0.6 \\ 81.32 \pm 0.6$	73.78 ± 0.3 67.32 ± 0.5 15.56 ± 0.6 83.79 ± 0.8 81.06 ± 0.6	72.05 ± 0.3 66.36 ± 1.0 13.69 ± 1.9 95.70 ± 1.8 93.23 ± 2.3	$\begin{array}{c} 44.83 \pm 0.0 \\ 32.69 \pm 0.3 \\ 20.75 \pm 0.5 \\ 97.41 \pm 0.2 \\ 91.92 \pm 0.2 \end{array}$	$68.76 \pm 0.1 66.02 \pm 0.4 4.70 \pm 0.7 92.69 \pm 0.3 91.71 \pm 0.2$	$74.04 \pm 0.1 \\ 72.18 \pm 0.0 \\ 4.47 \pm 0.0 \\ 92.76 \pm 0.8 \\ 74.53 \pm 0.8$	$83.71 \pm 0.2 \\ 80.54 \pm 0.2 \\ 7.62 \pm 0.4 \\ 84.81 \pm 0.9 \\ 71.52 \pm 0.5$
UTKFace	ACC Worst Gap EqOdd DP	$\begin{array}{c} 92.75 \pm 0.5 \\ 91.78 \pm 0.6 \\ 2.26 \pm 0.6 \\ 97.62 \pm 0.5 \\ 94.55 \pm 1.2 \end{array}$	$\begin{array}{c} 93.19 \pm 0.3 \\ 92.19 \pm 0.3 \\ 2.34 \pm 0.6 \\ 97.62 \pm 0.4 \\ 94.83 \pm 1.1 \end{array}$	$\begin{array}{c} 94.23 \pm 0.4 \\ 94.00 \pm 0.4 \\ 0.45 \pm 0.3 \\ 99.24 \pm 0.4 \\ 95.61 \pm 0.8 \end{array}$	$\begin{array}{c} 96.72 \pm 0.4 \\ 95.90 \pm 0.2 \\ 1.90 \pm 0.6 \\ 97.96 \pm 0.5 \\ 93.91 \pm 1.1 \end{array}$	$\begin{array}{c} 96.79 \pm 0.4 \\ 96.05 \pm 0.3 \\ 1.72 \pm 0.2 \\ 98.17 \pm 0.2 \\ 94.27 \pm 1.2 \end{array}$	$\begin{array}{c} 96.70 \pm 0.3 \\ 96.03 \pm 0.3 \\ 1.55 \pm 0.1 \\ 98.36 \pm 0.2 \\ 94.96 \pm 1.1 \end{array}$	$\begin{array}{c} 97.12 \pm 0.2 \\ 96.34 \pm 0.1 \\ 1.81 \pm 0.2 \\ 98.16 \pm 0.2 \\ 95.22 \pm 1.3 \end{array}$	$\begin{array}{c} 96.78 \pm 0.4 \\ 95.76 \pm 0.4 \\ 2.36 \pm 0.8 \\ 97.59 \pm 0.8 \\ 94.33 \pm 1.6 \end{array}$	$\begin{array}{c} 97.25 \pm 0.5 \\ 96.89 \pm 0.4 \\ 0.85 \pm 0.2 \\ 99.20 \pm 0.2 \\ 95.28 \pm 1.1 \end{array}$	$\begin{array}{c} 97.02 \pm 0.2 \\ 96.29 \pm 0.2 \\ 1.70 \pm 0.3 \\ 98.27 \pm 0.3 \\ 94.90 \pm 1.3 \end{array}$
FairFace	ACC Worst Gap EqOdd DP	$66.76 \pm 0.2 \\ 66.34 \pm 0.4 \\ 0.87 \pm 0.5 \\ 96.22 \pm 0.6 \\ 97.61 \pm 0.4$	68.37 ± 0.3 67.69 ± 0.3 1.44 ± 0.6 96.14 ± 0.1 97.55 ± 0.2	52.57 ± 0.0 50.74 ± 0.0 3.46 ± 0.0 93.86 ± 0.0 96.89 ± 0.0	$\begin{array}{c} 57.36 \pm 0.0 \\ 57.20 \pm 0.0 \\ 0.34 \pm 0.0 \\ 97.16 \pm 0.0 \\ 98.36 \pm 0.0 \end{array}$	$\begin{array}{c} 56.81 \pm 0.0 \\ 56.41 \pm 0.0 \\ 0.86 \pm 0.0 \\ 96.78 \pm 0.0 \\ 98.32 \pm 0.0 \end{array}$	$\begin{array}{c} 52.73 \pm 0.4 \\ 51.59 \pm 0.6 \\ 2.40 \pm 0.5 \\ 97.14 \pm 0.1 \\ 98.46 \pm 0.1 \end{array}$		65.57 ± 0.0 64.73 ± 0.0 1.79 ± 0.0 96.26 ± 0.0 97.57 ± 0.0	$\begin{array}{c} 59.40 \pm 0.1 \\ 56.71 \pm 0.0 \\ 5.71 \pm 0.1 \\ 95.13 \pm 0.1 \\ 97.16 \pm 0.1 \end{array}$	$\begin{array}{c} 56.52 \pm 0.2 \\ 53.53 \pm 0.0 \\ 6.35 \pm 0.4 \\ 96.54 \pm 0.1 \\ 98.00 \pm 0.1 \end{array}$
Facet	ACC Worst Gap EqOdd DP	67.55 ± 0.4 64.25 ± 1.0 4.31 ± 1.1 96.47 ± 1.2 95.40 ± 1.1	67.83 ± 0.7 64.94 ± 1.0 3.78 ± 1.0 96.68 ± 1.0 95.91 ± 1.0	41.16 ± 0.9 40.40 ± 0.9 3.22 ± 0.9 93.52 ± 1.4 94.10 ± 1.2	33.10 ± 0.2 31.43 ± 0.3 7.13 ± 0.9 98.72 ± 0.2 98.92 ± 0.3	33.17 ± 0.1 31.69 ± 0.2 6.32 ± 0.3 98.82 ± 0.3 99.00 ± 0.4	33.26 ± 0.2 31.54 ± 0.2 7.37 ± 1.2 99.75 ± 0.1 99.80 ± 0.1	$\begin{array}{c} 58.14 \pm 0.2 \\ 57.97 \pm 0.4 \\ 0.72 \pm 0.6 \\ 95.90 \pm 1.4 \\ 93.91 \pm 1.5 \end{array}$	68.41 ± 0.5 63.79 ± 1.3 6.04 ± 1.2 97.04 ± 1.2 96.15 ± 1.4	$\begin{array}{c} 57.75 \pm 0.7 \\ 57.25 \pm 0.5 \\ 1.42 \pm 1.2 \\ 96.22 \pm 1.0 \\ 94.56 \pm 0.6 \end{array}$	$67.54 \pm 0.4 \\ 64.61 \pm 1.0 \\ 3.83 \pm 1.0 \\ 97.55 \pm 1.1 \\ 96.28 \pm 1.0$
Waterbirds	ACC Worst Gap EqOdd DP	$85.63 \pm 1.4 \\ 84.20 \pm 0.9 \\ 2.87 \pm 0.9 \\ 66.53 \pm 3.3 \\ 77.67 \pm 3.5$	$86.09 \pm 1.4 \\ 84.52 \pm 2.6 \\ 3.14 \pm 2.5 \\ 68.99 \pm 2.5 \\ 77.25 \pm 3.0$	$\begin{array}{c} 52.30 \pm 0.0 \\ 39.18 \pm 0.0 \\ 26.23 \pm 0.0 \\ 68.24 \pm 0.0 \\ 63.48 \pm 0.0 \end{array}$	78.05 ± 0.0 74.53 ± 0.0 7.04 ± 0.0 73.96 ± 0.0 78.12 ± 0.0	77.77 ± 0.0 73.49 ± 0.0 8.56 ± 0.0 73.19 ± 0.0 76.73 ± 0.0	75.70 ± 0.3 72.81 ± 0.6 5.78 ± 0.5 95.57 ± 0.2 94.09 ± 0.6	$ \begin{vmatrix} 86.95 \pm 0.0 \\ 81.26 \pm 0.0 \\ 11.39 \pm 0.0 \\ 80.88 \pm 0.0 \\ 80.46 \pm 0.0 \end{vmatrix} $	$\begin{array}{c} 92.41 \pm 0.0 \\ 91.27 \pm 0.0 \\ 2.28 \pm 0.0 \\ 94.03 \pm 0.0 \\ 94.62 \pm 0.0 \end{array}$	$\begin{array}{c} 90.77 \pm 0.0 \\ 88.97 \pm 0.2 \\ 3.59 \pm 0.3 \\ 92.99 \pm 0.5 \\ 93.17 \pm 0.1 \end{array}$	$\begin{array}{c} 88.62 \pm 0.1 \\ 85.74 \pm 0.2 \\ 5.76 \pm 0.2 \\ 89.38 \pm 0.1 \\ 89.44 \pm 0.1 \end{array}$
HAM10000	AUC Worst Gap EqOdd DP	$88.35 \pm 1.8 \\ 84.68 \pm 2.0 \\ 4.11 \pm 2.1 \\ 88.17 \pm 3.1 \\ 82.22 \pm 4.8$	$89.09 \pm 1.3 \\ 84.67 \pm 3.3 \\ 4.99 \pm 3.7 \\ 88.43 \pm 4.0 \\ 84.73 \pm 3.4$	$\begin{array}{c} 38.87 \pm 3.0 \\ 39.00 \pm 5.0 \\ 6.89 \pm 4.4 \\ 98.19 \pm 1.1 \\ 97.89 \pm 1.1 \end{array}$	$\begin{array}{c} 52.15 \pm 2.6 \\ 51.56 \pm 4.0 \\ 4.14 \pm 3.9 \\ 96.24 \pm 2.5 \\ 99.09 \pm 0.6 \end{array}$	$\begin{array}{c} 51.88 \pm 2.9 \\ 52.01 \pm 3.9 \\ 3.12 \pm 2.7 \\ 95.75 \pm 2.4 \\ 98.94 \pm 0.7 \end{array}$	$52.53 \pm 3.9 \\ 53.41 \pm 4.6 \\ 2.24 \pm 2.1 \\ 96.19 \pm 3.0 \\ 98.32 \pm 0.5$	- - - -	- - -	- - -	- - -
Fitz17k	AUC Worst Gap EqOdd DP	$89.74 \pm 1.0 \\ 88.39 \pm 1.1 \\ 2.92 \pm 1.1 \\ 94.92 \pm 2.5 \\ 94.46 \pm 1.4$	$\begin{array}{c} 91.29 \pm 0.5 \\ 90.15 \pm 1.0 \\ 2.51 \pm 1.0 \\ 95.61 \pm 2.3 \\ 94.53 \pm 1.1 \end{array}$	67.08 ± 1.7 66.46 ± 1.9 3.74 ± 1.6 97.06 ± 1.6 98.40 ± 1.7	69.92 ± 0.8 69.78 ± 0.6 2.31 ± 1.5 89.87 ± 3.0 95.03 ± 2.0	$69.81 \pm 0.1 \\ 69.85 \pm 0.1 \\ 2.52 \pm 1.7 \\ 87.95 \pm 3.7 \\ 93.02 \pm 1.2$	69.37 ± 0.8 69.35 ± 1.7 4.73 ± 4.6 85.88 ± 6.2 92.19 ± 0.6	- - - -	- - -	- - - -	- - - -

F LIMITATIONS AND FUTURE DIRECTIONS

 Despite its breadth, NH-Fair is only an intermediate step toward a truly comprehensive fairness benchmark, yet several gaps remain:

- Dataset: Although we span three domains, most datasets are still for human faces, due to the scarcity of open, well-annotated fairness datasets. Some datasets include multiple protected attributes. Due to computational resource limitations, we did not fully use them but selected the attribute with a clear performance gap. Future work may consider intersectional groups (e.g., gender × race) and more visual domains. In addition, our zero-shot evaluations cannot guarantee that images were not used during the pretraining phase of LVLMs, which may lead to the potential data leakage problem.
- **Task scope**: To bridge single-modality bias mitigation algorithms with multi-modal models, we limited experiments to single-image classification. Segmentation, detection, and captioning could be further explored with task-specific fairness notions.
- Fairness Notions: As we stated in Section 2, there are also other fairness notions to quantify the model fairness, like individual fairness and counterfactual fairness, which could be used in future experiments.
- Experiments: We evaluated LVLMs in zero-shot only. Fine-tuning them could further reveal their fairness performance and bring more valuable insights. In the comparison, we divide the method into data-centric and algorithmic. However, combining all algorithmic approaches with data-centric approaches is a promising solution to achieve better fairness, which could be a future research direction.
- Analysis of LVLM Fairness: The paper considers fairness differences across LVLMs with different architectural or alignment choices but does not account for confounding factors like differences in training data or finetuning protocols due to the lack of transparency and standardization in recently released LVLMs. For example, the Gemma 3 Technical Report discloses only high-level information, such as the number of training tokens and the fact that both SFT and RLHF were used for alignment. Further details of implementation are unavailable, which prevents controlled comparisons. This lack of visibility is common in current LLM and LVLM releases, making some in-depth analysis difficult. As a result, our study intentionally focuses on aspects that are quantifiable and reproducible.

G BROADER IMPACTS

NH-Fair provides a transparent and public benchmark that lets researchers and developers evaluate both single-modal and vision-language models under the same fairness metrics, motivating the community to build more fair machine learning models. However, the visibility also may lead to "benchmark gaming": a model can be fine-tuned to excel on the reported sub-groups while quietly marginalizing untested or intersectional minorities, posing potential downstream AI-safety concerns.

H LLM USAGE

In accordance with the ICLR policy, we disclose the use of large language models (LLMs) in preparing this paper. LLMs are also a part of the subjects of our research, and all experimental results are based on our own implementations and evaluations of these models. Separately, we used LLMs as general-purpose assistive tools for language editing, improving the clarity of writing. They were not used for research ideation, system design, or the generation or analysis of experimental results.