

Mitigating Bias with Words: Inducing Demographic Ambiguity in Face Recognition Templates by Text Encoding

Tahar Chettaoui¹

tahar.chettaoui@igd.fraunhofer.de

Naser Damer^{1,2}

naser.damer@igd.fraunhofer.de

Fadi Boutros¹

fadi.boutros@igd.fraunhofer.de

¹ Fraunhofer IGD

Fraunhoferstraße 5

64283 Darmstadt

² TU Darmstadt

Karolinenpl. 5,

64289 Darmstadt

Abstract

Face recognition (FR) systems are often prone to demographic biases, partially due to the entanglement of demographic-specific information with identity-relevant features in facial embeddings. This bias is extremely critical in large multicultural cities, especially where biometrics play a major role in smart city infrastructure. The entanglement can cause demographic attributes to overshadow identity cues in the embedding space, resulting in disparities in verification performance across different demographic groups. To address this issue, we propose a novel strategy, Unified Text-Image Embedding (UTIE), which aims to induce demographic ambiguity in face embeddings by enriching them with information related to other demographic groups. This encourages face embeddings to emphasize identity-relevant features and thus promotes fairer verification performance across groups. UTIE leverages the zero-shot capabilities and cross-modal semantic alignment of Vision-Language Models (VLMs). Given that VLMs are naturally trained to align visual and textual representations, we enrich the facial embeddings of each demographic group with text-derived demographic features extracted from other demographic groups. This encourages a more neutral representation in terms of demographic attributes. We evaluate UTIE using three VLMs, CLIP, OpenCLIP, and SigLIP, on two widely used benchmarks, RFW and BFW, designed to assess bias in FR. Experimental results show that UTIE consistently reduces bias metrics while maintaining, or even improving in several cases, the face verification accuracy.

1 Introduction

Face recognition (FR) systems are increasingly integrated into real-world applications [16, 17, 18], yet concerns around demographic bias remain a barrier to their fair and reliable deployment [8, 19, 20, 21]. Such biases manifest as disparities in recognition accuracy across demographic groups, including race, gender, and age [1, 2, 16, 19, 20]. Addressing these disparities is essential to ensure equitable FR performance, motivating the development of methods [3, 22, 23, 24, 25] to mitigate bias while maintaining high recognition accuracy. Despite progress in developing bias mitigation techniques, effectively reducing disparities across demographic groups in verification scenarios remains challenging. This challenge arises partly because embeddings used for face verification may encode demographic-specific attributes

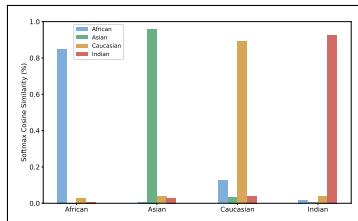
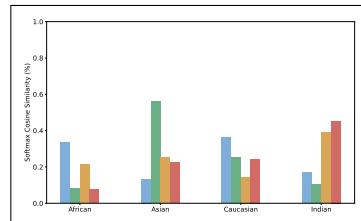
(a) Baseline (IE), with the template defined as T_i-hat (b) UTIE, with the template defined as T_i-hat

Figure 1: Comparison of our UTIE average cosine similarities with racial text embeddings against the baseline (IE) on the RFW subsets, using CLIP ViT/B-16. For each subset, we compute the average cosine similarities of the face embeddings with the different text embeddings T_i , for $i \in 1, \dots, 4$, representing the four considered races: African (blue), Asian (green), Caucasian (yellow), and Indian (red). (a) IE displays a high similarity with the corresponding demographic class T_i-hat for all subsets while maintaining lower similarities with the other classes T_i for $i \in 1, \dots, 4$, $i \neq \hat{i}$. This shows that face embeddings encode demographic-specific information. (b) For UTIE, the similarity to T_i-hat is significantly lower compared to IE, and the similarities are spread more evenly across the other classes. This indicates that UTIE makes the racial information in the embeddings more ambiguous, reducing the alignment with a single demographic and distributing similarity more evenly across classes.

that contribute to biased performance across groups [8, 10]. In this work, we address this issue, focusing on disparities in verification performance across demographic groups.

Since verification relies on comparing embeddings extracted from face images, we hypothesize that one contributing factor to this bias is the presence of demographic-specific information in these embeddings. To mitigate this, we propose a novel approach that encourages face embeddings to be more demographically ambiguous by enriching them with features representative of all demographic groups. This strategy aims to make the learned representations more generalizable and less sensitive to demographic attributes, thereby reducing bias in face verification performance. To achieve this, we leverage multimodal Vision-Language Models (VLMs) and the fact that they produce aligned embeddings for image and text. We use the text encoder to obtain embeddings that capture demographic characteristics. Next, we exclude the embedding corresponding to the predicted characteristic (e.g., racial, gender), i.e., the textual embedding that matches the image embedding with the highest similarity score. We then compute the average of the remaining embeddings and add this mean vector to the original image embedding. By incorporating features representative of other demographic groups into the final visual feature representation, we aim to make the influence of demographic characteristics more ambiguous within the face embedding. The evaluations include demographic bias analysis on the Racial Faces in the Wild (RFW) [32] and Balanced Faces In the Wild (BFW) [25] datasets, reporting verification accuracies along with bias assessment metrics, including the Skewed Error Ratio (SER) and standard deviation (STD). Both datasets evaluate racial bias, and BFW additionally includes gender bias evaluation. Furthermore, we evaluate the demographic prediction performance in a zero-shot setting on the same datasets since our approach relies on accurate and consistent demographic attribute predictions from the text encoder to construct the enriched face embeddings. We considered three VLMs, namely CLIP [27], OpenCLIP [5], and SigLIP [44], and our results reveal that our unified text-image embedding (UTIE) successfully reduces the STD and SER while improving or maintaining comparable mean accuracy across all models on both racial and

gender bias evaluation. For instance, in the racial bias evaluation, baseline CLIP achieves a mean accuracy of 72.20% on RFW, which increases to 72.25% using UTIE, while the STD decreases using our UTIE from 4.81 to 4.46 and the SER from 1.50 to 1.45. Similarly, in the gender bias evaluation on BFW, using UTIE with CLIP reduces the STD from 2.72 to 2.58 and lowers the SER from 1.28 to 1.27, while maintaining comparable accuracy. These results confirm that UTIE enhances demographic ambiguity in the embeddings, leading to reduced bias while preserving or improving recognition performance. These findings contribute to ongoing research on bias mitigation in FR and support the development of demographically ambiguous face embeddings as a viable strategy for reducing bias.

2 Related Work

Demographic bias in FR has received significant attention. Several works [11, 12, 13] examined the causes of demographic bias, identifying imbalanced or unrepresentative datasets as a factor. Albiero et al. [11] identified the under-representation of females in training data as a key factor contributing to gender bias. Another study [12] showed that training FR systems on balanced datasets is critical for reducing vulnerabilities on specific demographic groups. Another factor is the variability in skin tone [13, 14]. Buolamwini and Gebru [13] evaluated three gender classification systems and identified under-performance in recognizing dark-skinned females. A quantitative assessment across six skin tone groups found light-skinned individuals easiest to verify and darker-skinned individuals the most challenging [14]. Other factors, like the impact of image attributes associated with socially constructed gender norms, have also been investigated. Prior works [15, 16, 17] explored the impact of makeup on FR accuracy, finding that variations in makeup between images make it more challenging to recognize genuine pairs. Other studies [18, 19] examined facial hair, showing that similar beard areas increase similarity scores, while differences reduce them.

All these papers motivated research into how to mitigate bias by modifying the training strategy or the network structure. To promote fairness in FR, [20] proposes a reinforcement learning-based Race Balance Network (RL-RBN) that uses adaptive margins to balance FR performance across races. By modeling margin selection for non-Caucasians as a Markov decision process and applying deep Q-learning, RL-RBN reduces feature distribution skewness. Alvi et al. [21] propose an algorithm that removes multiple sources of variation from the feature representation of a network to reduce bias and improve classification accuracy. They introduce a supervised learning approach that learns features informative for the primary task while remaining uninformative for spurious variations such as pose or gender, thereby reducing bias in the learned embeddings. Dhar et al. [22] proposed Protected Attribute Suppression System (PASS), an adversarial debiasing approach to train a model to classify identities while discouraging it from predicting protected attributes. It can be trained on descriptors from any trained high-performing network to classify identities while reducing the encoding of sensitive attributes, effectively mitigating gender and skin tone bias. Another work [23] proposed a meta-learning algorithm, Meta Balanced Network (MBN), to mitigate algorithmic bias, which learns adaptive margins in large margin loss. MBN balances the feature scatter among skin tone groups in feature space rather than merely adjusting loss weights.

None of the previous works addressed mitigating bias in VLMs, especially those that are not fine-tuned to the task, which we address in this work. Additionally, and most importantly, this paper is the first to theorize reducing demographic bias in FR by making the face embeddings more demographically ambiguous, and doing that by inducing embeddings encoded from text to the face image embedding. Unlike prior works that encourage balance during training [20, 23], discourage the encoding of demographic attributes in the feature

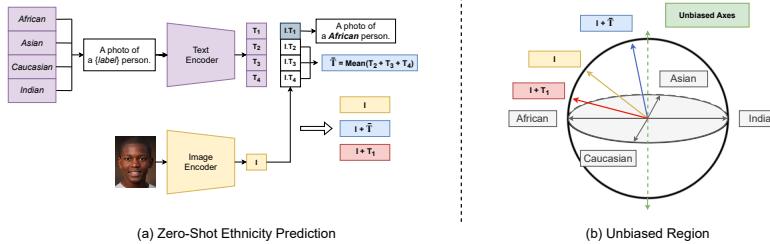


Figure 2: Demographically ambiguous embedding generation with UTIE for racial bias mitigation. (a) We use a text encoder $f_t(\cdot)$ to extract racial embeddings T_i , for $i \in 1, \dots, 4$ for zero-shot prediction. We then exclude the predicted embedding T_i (in this case T_1), compute the average of the remaining demographic embeddings, denoted as \bar{T} , and add it to the original image embedding I resulting to I' . (b) The yellow vector represents the baseline IE. The blue vector represents our UTIE. This vector is closer to the unbiased axes, as it is more demographically ambiguous, as demonstrated in Section 3.2. Finally, the red vector represents IE+PTE defined in Section 4. This vector shows stronger demographic dominance, in this case leaning more toward African identity, illustrating increased demographic influence.

representation [2, 3], or manipulate the comparison and decision-making processes [29, 30], our approach enriches each face embedding with information derived from embeddings that explicitly capture characteristics of other demographic groups. This balances the influence of demographic attributes and reduces their dominance within the embeddings, resulting in a more demographically ambiguous face representation and thus reducing demographic bias.

3 Methodology

This section introduces our approach to reducing demographic bias in VLM-based FR by making face embeddings more demographically ambiguous. We first present baseline VLMs, which learn to align images and text embeddings, allowing us to leverage complementary semantic information from text. We then introduce UTIE, showing how we fuse image and text embeddings to obscure demographic information, aiming at reducing demographic bias.

3.1 Baseline VLMs (preliminary)

In this paper, we considered three VLMs, namely CLIP [2], OpenCLIP [3], and SigLIP [4]. Some of which have been shown to be usable as template extractors for FR [5]. These methods are based on contrastive learning, which aligns related image-text pairs by minimizing the distance between their visual and textual embeddings in a shared space, while maximizing the distance between embeddings of unrelated pairs [6, 13, 20, 21]. Given a batch of N (image, text) pairs, the model is trained to identify which of the $N \times N$ possible pairings correspond to the true pairs. To achieve this, they jointly train an image encoder $f_i(\cdot)$ and a text encoder $f_t(\cdot)$ to align the image and text embeddings for N matching (positive) image-text pairs while making sure that unrelated $N_2 - N$ (negative) image-text pairs are dissimilar in the embedding space. This is done by maximizing the cosine similarity of the image and text embeddings of the matching pairs in the batch while minimizing the cosine similarity of the embeddings of the incorrect pairings. SigLIP is trained using a sigmoid loss instead of the softmax-based contrastive loss. It operates solely on individual image-text pairs and does not require a global view of pairwise similarities for normalization. This allows for scaling to larger batch sizes and improves performance at smaller batch sizes. The nature of these pre-trained VLMs allows the extraction of embeddings from face images

that can be used for face verification [40]. They also enable generating embeddings from text that are aligned with those of the face images. This alignment between text and image embeddings forms the basis of our UTIE solution, as described in the next Section 3.2.

3.2 Unified Text-Image Embedding (UTIE)

This section presents our UTIE, which combines image and text embeddings to obscure demographic-specific information and reduce demographic bias. Since face verification compares embeddings that may encode demographic cues, we propose an approach that encourages face embeddings to be more demographically ambiguous by enriching them with features representative T_i of all N demographic groups. Specifically, we enrich each face embedding with information derived from embeddings that are explicitly designed to capture characteristics of other (not dominant in this face) demographic groups. To achieve this, we leverage VLMs, which align image and text embeddings in the same space, allowing us to enrich face embeddings with demographic features, as detailed in Section 3.1.

As depicted in Figure 2.a, given a sample face image x , we pass it through the image encoder to obtain its image embedding I . Next, we aim to generate features that are representative of all demographic groups. To achieve this, we build on the alignment within the image-text embedding space and propose leveraging the text encoder to generate these representative features. Following the zero-shot settings described in Section 4, we extract text embeddings, denoted as T_i , for the N class labels corresponding to different demographic characteristics. Since the prompt specifies only the demographic classes, the text embeddings T_i , for $i \in 1, \dots, N$, capture information that is ideally exclusively related to those demographics. We then use these embeddings for zero-shot prediction, where we compute the predicted demographic class as the one with the maximum cosine similarity to I :

$$\hat{i} = \arg \max_i \frac{\langle I, T_i \rangle}{\|I\| \cdot \|T_i\|}, \quad (1)$$

As a result, we obtain the index \hat{i} , corresponding to the predicted demographic class. This index indicates the demographic group that the model predicts for I based on the maximum alignment in the embedding space. To encourage demographic ambiguity, we aim to shift the image embedding away from this dominant demographic group. To do that, we need to enrich the face image embedding with information about the other classes by constructing the mean text embedding across all demographic groups, excluding the predicted one:

$$\bar{T} = \frac{1}{N-1} \sum_{i \neq \hat{i}} T_i, \quad (2)$$

where N is the number of demographic classes. We then add \bar{T} to the image embedding to obtain the final demographically ambiguous representation using our UTIE:

$$I' = I + \bar{T}. \quad (3)$$

In our UTIE solution, we propose to use this adjusted embedding I' as the face template for verification instead of the original embedding I . By adding the mean embedding \bar{T} to the image embedding I , this operation aims to reduce the dominance of the predicted demographic class $T_{\hat{i}}$. As a result, I' distributes more evenly across the directions associated with all demographic groups. This makes the resulting feature vector I' less aligned with any single demographic group and more balanced across all groups, making it demographically ambiguous and helping to reduce demographic bias in downstream tasks.

Figure 1 compares the cosine similarity distributions of image embeddings with racial text embeddings across the four demographic classes on the RFW [34], introduced in the next section, subsets for both the baseline (IE) and our proposed approach (UTIE). For all subsets, IE displays a high similarity with the corresponding demographic class $T_{\hat{i}}$ while maintaining notably lower similarities with the other classes T_i for $i \in 1, \dots, N$, $i \neq \hat{i}$. On the other hand,

for our approach (UTIE), the similarity to $T_{\hat{i}}$ is significantly lower compared to IE, and the similarities spread more evenly across other classes. This indicates that our method makes the racial information in the embeddings more ambiguous, reducing the strong alignment with a single demographic and distributing similarity more evenly across classes.

4 Experimental Setup

Face Template Setups. To validate our approach, we consider three feature representations. The first configuration uses the unmodified image embedding (IE) as a baseline, with the template defined as I in Section 3.2. The second configuration applies our proposed UTIE approach, with the template defined as I' in Section 3.2. This aims to mitigate bias, making demographic influences more ambiguous in the representation. As a third configuration, to demonstrate that performing the counter concept to UTIE increases bias, we introduce the IE+PTE setup, where PTE stands for Predicted demographic Text Embedding. In this setup, instead of aiming at reducing demographic information, we attempt to reinforce it by adding the predicted demographic text embedding $T_{\hat{i}}$ to the image embedding I , resulting in:

$$I^* = I + T_{\hat{i}}. \quad (4)$$

This configuration allows us to examine whether explicitly incorporating the predicted demographic attribute amplifies demographic bias in the embeddings, providing a contrasting perspective that highlights the effectiveness of our UTIE approach in reducing bias.

Evaluation Datasets. We evaluate the considered VLMs on the RFW [34] and BFW [23] datasets to assess model bias and performance across demographic groups, as they are widely used for fairness and bias evaluation [8, 25]. The RFW dataset contains four testing subsets corresponding to African, Asian, Caucasian, and Indian groups. The BFW dataset extends this evaluation by providing eight subgroups balanced across gender and race, enabling a more detailed assessment of intersectional bias. For racial bias evaluation, both RFW and BFW are suitable, while gender bias is assessed only with BFW.

Zero-Shot Settings. Following the evaluation benchmarks [25, 34] for racial bias evaluation, we consider four ethnicities as classes: African, Asian, Caucasian, and Indian. CLIP [2], OpenCLIP [5], and SigLIP [40] were trained using sentence-based prompts rather than single-word labels. This motivates using prompts during zero-shot evaluation, such as "A photo of a {label} .", to help bridge the distribution gap between pre-training and downstream tasks, which has been shown to improve accuracy. Additionally, it was observed that zero-shot performance can be significantly enhanced by customizing the prompt text for each task [22]. For these reasons, we adopt the prompt "A photo of a {label} person ." as our standard prompt in this work. Following the same setup, we evaluate gender bias in Section 5.2 using two classes, female and male, instead of four ethnic groups.

Evaluation Metrics. We follow the reporting protocols and evaluation metrics associated with the evaluation datasets and recent works [8, 24, 25]. We report the results as verification accuracies in (%) on each subset and as average accuracies to evaluate general recognition performance on the benchmarks. To evaluate the bias, we report the STD between all subsets and the SER, which is given by $\frac{\max_g Error_g}{\min_g Error_g}$, where g represents the demographic group, as reported in [25, 23]. A higher STD value indicates more bias across demographic groups and vice versa. For SER, models with values closer to 1 are less biased.

Model Architectures. We use official pre-trained models for all three VLMs: CLIP (ViT-B/16) [2], OpenCLIP (ViT-B/16) [5], and SigLIP (ViT-B/16) [40]. The OpenCLIP model was trained on the LAION-2B dataset [2], which includes a 2.32 billion-sample English image-text subset. For comparison, CLIP [2] was trained on a smaller dataset of

400 million (image, text) pairs sourced from large-scale publicly available data. For SigLIP, we use the ViT-B/16 model trained on the WebLI dataset [4] at a resolution of 224x224.

5 Results

This section presents the results achieved by our UTIE solution on the considered VLMs on RFW and BFW. We evaluate three feature representation strategies, IE, UTIE, and IE+PTE, presented in Section 4. We assess these configurations for racial bias in Section 5.1 and gender bias in Section 5.2, examining their impact on fairness in verification performance.

5.1 Racial Bias in FR Evaluation

We evaluated the considered approaches, namely CLIP [2], OpenCLIP [5], and SigLIP [40], on the RFW [32] and BFW [25] benchmarks to assess model bias and performance across different demographic groups. To evaluate our proposed UTIE, we compare it with the baseline (IE) and a validation setup (IE+PTE) where we add the predicted demographic embedding to amplify bias signals. This setup allows us to systematically evaluate the impact of each feature representation on racial bias. One can observe the following from the results in Table 5.1, highlighting the impact of different feature representations on racial bias:

Comparison of UTIE vs. IE: UTIE reduces STD and SER across all settings on both datasets, while improving mean accuracy on RFW [32] and maintaining comparable accuracy on BFW [25]. For instance, CLIP [2] achieves a mean accuracy of 72.20% on RFW with IE and 72.25% with UTIE, while reducing STD from 4.81 to 4.46 and SER from 1.50 to 1.45. On BFW, CLIP with UTIE maintains a high mean accuracy (84.23% vs. 84.50%) while reducing STD (1.54 vs. 1.63) and SER (1.26 vs. 1.29). This trend holds for OpenCLIP [5] and SigLIP [40], demonstrating that UTIE, by making face embeddings more demographically ambiguous as shown in Figure 1, can preserve performance and reduce bias. In contrast, the IE+PTE setup yields higher STD and SER, confirming that adding predicted demographic embeddings amplifies demographic bias as intended for validation.

Comparison across models: CLIP [2] and OpenCLIP [5] consistently achieve higher average accuracy across both RFW [32] and BFW [25] compared to SigLIP [40], with BFW means around 84% for both CLIP and OpenCLIP, while SigLIP reaches around 79%. SigLIP exhibits lower STD and SER on BFW, reaching an STD of 0.97 with UTIE compared to 1.54 and 2.13 for CLIP and OpenCLIP, respectively, indicating stronger consistency across demographic groups despite its lower overall accuracy. Compared to OpenCLIP, CLIP is more balanced across demographic groups, exhibiting lower STD and SER in all cases.

Comparison of Dataset influence on accuracy and bias: Across all models, the average accuracy is consistently higher on BFW [25] than on RFW [32], consistent with previous works [40, 12, 25]. For example, baseline CLIP [2] achieves an average accuracy of 84.50% on BFW compared to 72.20% on RFW. Additionally, both the STD and SER are lower on BFW across all models, indicating reduced demographic variance compared to RFW.

Overall, our UTIE consistently reduced the bias of the investigated model over racial groups, while having minimal effect on the FR performance. This, along with the clear negative effect of the counter concept, IE+PTE, validates our approach of reducing demographic bias by making the demographic information in the embedding more ambiguous.

5.2 Gender Bias Evaluation

Besides racial bias, we assess gender bias using the BFW dataset [25]. We follow the same evaluation settings used for the racial bias assessment and examine the same approaches, focusing on their performance across the two gender groups. From Table 5.2, highlighting the impact of different feature representations on gender bias, the following can be noted:

Approach	Feature Embedding	RFW [32]							BFW [25]						
		African	Asian	Caucasian	Indian	Mean	STD	SER	Asian	Black	Indian	White	Mean	STD	SER
CLIP [2]	IE	70.75	69.73	79.32	68.98	72.20	4.81	1.50	82.36	84.49	84.85	86.31	84.50	1.63	1.29
	UTIE	70.85	69.80	78.88	69.48	72.25	4.46	1.45	82.20	84.16	84.68	85.89	84.23	1.54	1.26
	IE+PTE	68.47	68.73	77.90	66.87	70.49	5.01	1.50	82.22	83.37	83.37	86.25	83.95	1.70	1.29
OpenCLIP [8]	IE	69.37	68.60	79.95	69.72	71.91	5.38	1.57	80.49	86.01	83.97	86.35	84.20	2.69	1.43
	UTIE	69.35	68.83	79.80	69.85	71.96	5.24	1.54	80.54	85.23	83.79	84.85	83.60	2.13	1.32
	IE+PTE	67.83	67.62	78.93	65.58	69.99	6.05	1.63	80.07	84.85	82.63	86.37	83.48	2.74	1.46
SigLIP [40]	IE	58.17	64.17	71.62	65.98	64.98	5.54	1.47	78.27	79.52	80.57	80.54	79.73	1.09	1.12
	UTIE	58.63	64.62	71.63	65.60	65.12	5.32	1.46	77.83	78.91	79.93	79.80	79.12	0.97	1.10
	IE+PTE	56.80	61.02	69.86	63.53	62.80	5.46	1.43	77.77	79.17	79.93	80.88	79.44	1.31	1.16

Table 1: Evaluation results of different feature representations for the considered models on RFW [32] and BFW [25] reported as average verification accuracy in (%), STD, and SER across four different demographic groups. The higher mean indicates better verification accuracy, and the higher STD indicates a more biased model. For SER, the model that achieved a value closer to 1 is less biased. UTIE consistently reduces STD and SER across all settings on both datasets, while maintaining comparable accuracy on RFW and BFW.

Comparison of UTIE vs. IE: For all models, UTIE shows only a slight decrease in average recognition performance compared to baseline IE. Additionally, UTIE consistently reduces both STD and SER across the two gender groups on BFW, demonstrating that increasing demographic ambiguity in embeddings can effectively reduce gender bias. For example, with UTIE, OpenCLIP [8] improves, with STD dropping from 3.86 to 3.06 and SER from 1.42 to 1.30. CLIP [2] STD decreases from 2.72 to 2.58, with a slight reduction in SER from 1.28 to 1.27. SigLIP [40] maintains a stable STD while showing a small SER improvement. The IE+PTE configuration consistently demonstrates lower accuracy as well as higher STD and SER, confirming that incorporating predicted demographic text embeddings amplifies demographic influence and bias in representations, as expected for validation. These observations follow the trend observed in the racial bias evaluation in Section 5.1.

Comparison across models: CLIP [2] generally achieves the highest mean accuracy, followed by OpenCLIP [8], with SigLIP [40] showing lower performance. OpenCLIP reports a higher STD of 3.06 and SER of 1.30 than CLIP, which achieves an STD of 2.58 and SER of 1.27, indicating that OpenCLIP is slightly more biased across genders, although UTIE helps reduce this gap. SigLIP consistently yields lower mean accuracy but demonstrates the lowest STD of 1.59 and SER of 1.11, reflecting lower gender bias.

Our UTIE reduces gender bias in the evaluated model without significantly impacting FR accuracy, similar to the trend observed in the racial bias evaluation in Section 5.1.

Approach	Feature Embedding	Female	Male	Mean	STD	SER
CLIP [2]	IE	82.58	86.43	84.50	2.72	1.28
	UTIE	82.58	86.23	84.41	2.58	1.27
	IE+PTE	82.65	86.58	84.61	2.78	1.29
OpenCLIP [8]	IE	81.48	86.93	84.20	3.86	1.42
	UTIE	81.44	85.76	83.60	3.06	1.30
	IE+PTE	81.25	86.74	83.99	3.88	1.41
SigLIP [40]	IE	78.60	80.85	79.73	1.59	1.12
	UTIE	78.13	80.38	79.25	1.59	1.11
	IE+PTE	78.29	80.76	79.52	1.75	1.13

Table 2: Evaluation results of the different feature representations for the considered models on BFW [25] reported as average verification accuracy in (%), STD, and SER across the two gender groups. The UTIE consistently reduces STD and SER across all settings.

5.3 Zero-shot text encoder performance:

Our approach relies on consistent demographic attribute predictions from the text encoder to construct the enriched face embeddings. This is based on the need to determine the demographic group \hat{i} for the image embedding I , as used in Equation 1. Therefore, following the zero-shot settings detailed in Section 4, we evaluate the zero-shot prediction accuracy of the

text encoder for both race and gender. From Table 5.3, which presents the zero-shot accuracy of the text encoder for racial and gender prediction, we draw the following observations:

Analysis of zero-shot text encoder performance in recognizing racial groups: All models achieve high mean accuracy across groups, with CLIP [27] and OpenCLIP [5] reaching around 92.93% on RFW [52] and 87.92% and 91.44%, respectively, on BFW [52]. In comparison, SigLIP [40] shows a lower mean accuracy, scoring 88.07% on RFW and 86.15% on BFW, reflecting its overall lower zero-shot text encoder performance.

Analysis of zero-shot text encoder performance in recognizing gender groups: All models achieve high mean accuracy across gender groups, with CLIP [27] and OpenCLIP [5] reaching 98.42% and 97.42% on BFW, respectively, while SigLIP [40] shows lower performance at 95.11%, highlighting its generally lower zero-shot text encoder accuracy.

Approach	Architecture	RFW [52]				BFW [52]								
		African	Asian	Caucasian	Indian	Mean	Asian	Black	Indian	White	Mean	Female	Male	Mean
CLIP [27]	ViT-B/16	87.72	96.34	97.31	89.74	92.78	98.80	64.00	90.84	98.04	87.92	97.86	98.98	98.42
OpenCLIP [5]	ViT-B/16	95.99	93.74	97.44	84.57	92.94	95.16	85.40	85.88	99.32	91.44	96.78	98.06	97.42
SigLIP [40]	ViT-B/16	92.68	82.08	96.77	80.75	88.07	83.68	79.20	82.82	98.90	86.15	92.68	97.54	95.11

Table 3: Zero-shot racial prediction accuracy on RFW [52] and BFW [52], and gender prediction accuracy on BFW, for the considered models, reported as average prediction accuracy (%). OpenCLIP [5] achieves the highest mean zero-shot accuracy for racial prediction on RFW and BFW, followed by CLIP [27], while SigLIP [40] shows the lowest mean accuracy. For gender prediction on BFW, CLIP achieves the highest mean zero-shot accuracy.

Even though performing this recognition is neither our contribution nor our goal, to put the reported results in a wider perspective, the performance compares well to recent works on the same benchmarks, with [39] reporting 93.66% average face verification accuracy.

6 Conclusion

We proposed a novel approach to reduce demographic bias in VLM-based FR by inducing demographic ambiguity in face embeddings. We employed VLMs to create a UTIE that incorporates features from multiple demographic groups to introduce demographic ambiguity, thereby reducing the dominance of any single demographic attribute in the face representation while preserving identity information. We experimentally demonstrated this ambiguity by showing a more evenly distributed demographic similarity distribution on demographic groups after applying our approach, compared to the baseline. Our approach was evaluated on RFW and BFW, using three SOTA models: CLIP, OpenCLIP, and SigLIP. Our results consistently demonstrate a reduction in bias metrics while maintaining verification accuracy. For example, when comparing the baseline with UTIE, we observe that UTIE reduces the STD from 4.81 to 4.46 and the SER from 1.50 to 1.45 on RFW with CLIP while improving the mean accuracy from 72.20% to 72.25%. On BFW, CLIP with UTIE maintains a comparable mean accuracy while reducing the STD from 1.63 to 1.54 and the SER from 1.29 to 1.26. This trend holds across OpenCLIP and SigLIP. These findings validate that increasing demographic ambiguity in face embeddings is a promising direction for reducing bias in FR systems and motivate further research on improving this strategy, including prompt engineering, ensembling methods, and investigating its potential impact on privacy leakage.

Acknowledgments

This research work has been funded by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE.

References

- [1] Vítor Albiero, Krishnapriya K. S, Kushal Vangara, Kai Zhang, Michael C. King, and Kevin W. Bowyer. Analysis of gender inequality in face recognition accuracy. In *WACV Workshops*, pages 81–89. IEEE, 2020.
- [2] Mohsan S. Alvi, Andrew Zisserman, and Christoffer Nellåker. Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In *ECCV Workshops (1)*, volume 11129 of *Lecture Notes in Computer Science*, pages 556–572. Springer, 2018.
- [3] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *FAT*, volume 81 of *Proceedings of Machine Learning Research*, pages 77–91. PMLR, 2018.
- [4] Xi Chen, Xiao Wang, Soravit Changpinyo, A. J. Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, Alexander Kolesnikov, Joan Puigcerver, Nan Ding, Keran Rong, Hassan Akbari, Gaurav Mishra, Linting Xue, Ashish V. Thapliyal, James Bradbury, and Weicheng Kuo. Pali: A jointly-scaled multilingual language-image model. In *ICLR*. OpenReview.net, 2023.
- [5] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *CVPR*, pages 2818–2829. IEEE, 2023.
- [6] Tahar Chettaoui, Naser Damer, and Fadi Boutros. Froundation: Are foundation models ready for face recognition? *Image Vis. Comput.*, 156:105453, 2025.
- [7] Antitza Dantcheva, Cunjian Chen, and Arun Ross. Can facial cosmetics affect the matching accuracy of face recognition systems? In *BTAS*, pages 391–398. IEEE, 2012.
- [8] Prithviraj Dhar, Joshua Gleason, Aniket Roy, Carlos Domingo Castillo, and Rama Chellappa. PASS: protected attribute suppression system for mitigating bias in face recognition. In *ICCV*, pages 15067–15076. IEEE, 2021.
- [9] Paweł Drozdowski, Christian Rathgeb, Antitza Dantcheva, Naser Damer, and Christoph Busch. Demographic bias in biometrics: A survey on an emerging challenge. *CoRR*, abs/2003.02488, 2020.
- [10] Alexandre Fournier-Montgieux, Michaël Soumm, Adrian Popescu, Bertrand Luvison, and Hervé Le Borgne. Fairer analysis and demographically balanced face generation for fairer face verification. In *WACV*, pages 2788–2798. IEEE, 2025.
- [11] Patrick Grother, Mei Ngan, and Kayee Hanaoka. Face recognition vendor test part 3: Demographic effects, 2019.
- [12] Guodong Guo, Lingyun Wen, and Shuicheng Yan. Face authentication with makeup changes. *IEEE Trans. Circuits Syst. Video Technol.*, 24(5):814–825, 2014.
- [13] Marco Huber, Anh Thi Luu, Fadi Boutros, Arjan Kuiper, and Naser Damer. Bias and diversity in synthetic-based face recognition. In *WACV*, pages 6203–6214. IEEE, 2024.

- [14] Anil K. Jain, Debayan Deb, and Joshua J. Engelsma. Biometrics: Trust, but verify. *IEEE Trans. Biom. Behav. Identity Sci.*, 4(3):303–323, 2022.
- [15] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 4904–4916. PMLR, 2021.
- [16] Christopher Jones. Law enforcement use of facial recognition: Bias, disparate impacts to people of color, and the need for federal legislation. *N.C. J.L. and Tech.*, 22:777, 2021.
- [17] Minchul Kim, Anil Jain, and Xiaoming Liu. 50 years of automated face recognition. *CoRR*, abs/2505.24247, 2025.
- [18] Brendan Klare, Mark James Burge, Joshua C. Klontz, Richard W. Vorder Bruegge, and Anil K. Jain. Face recognition performance: Role of demographic information. *IEEE Trans. Inf. Forensics Secur.*, 7(6):1789–1801, 2012.
- [19] Ketan Kotwal and Sébastien Marcel. Review of demographic bias in face recognition. *CoRR*, abs/2502.02309, 2025.
- [20] Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, pages 12888–12900. PMLR, 2022.
- [21] Boyu Lu, Jun-Cheng Chen, Carlos Domingo Castillo, and Rama Chellappa. An experimental evaluation of covariates effects on unconstrained face verification. *IEEE Trans. Biom. Behav. Identity Sci.*, 1(1):42–55, 2019.
- [22] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021.
- [23] Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, and Emily Denton. Saving face: Investigating the ethical concerns of facial recognition auditing. In *AIES*, pages 145–151. ACM, 2020.
- [24] Christian Rathgeb, Paweł Drozdowski, Dinusha C. Frings, Naser Damer, and Christoph Busch. Demographic fairness in biometric systems: What do the experts say? *IEEE Technol. Soc. Mag.*, 41(4):71–82, 2022.
- [25] Joseph P. Robinson, Gennady Livitz, Yann Henon, Can Qin, Yun Fu, and Samson Timoner. Face recognition: Too bias, or not too bias? In *CVPR Workshops*, pages 1–10. Computer Vision Foundation / IEEE, 2020.
- [26] Krishnapriya K. S, Kushal Vangara, Michael C. King, Vitor Albiero, and Kevin W. Bowyer. Characterizing the variability in face recognition accuracy relative to race. In *CVPR Workshops*, pages 2278–2285. Computer Vision Foundation / IEEE, 2019.

[27] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5B: an open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022.

[28] Ignacio Serna, Aytami Morales, Julian Fiérrez, Manuel Cebrián, Nick Obradovich, and Iyad Rahwan. Sensitiveloss: Improving accuracy and fairness of face representations with discrimination-aware deep learning. *CoRR*, abs/2004.11246, 2020.

[29] Philipp Terhörst, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Post-comparison mitigation of demographic bias in face recognition using fair score normalization. *Pattern Recognit. Lett.*, 140:332–338, 2020. doi: 10.1016/j.patrec.2020.11.007. URL <https://doi.org/10.1016/j.patrec.2020.11.007>.

[30] Philipp Terhörst, Mai Ly Tran, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Comparison-level mitigation of ethnic bias in face recognition. In *IWBF*, pages 1–6. IEEE, 2020.

[31] Sayako Ueda and Takamasa Koyama. Influence of make-up on facial recognition. *Perception*, 39:260–4, 02 2010. doi: 10.1080/p6634.

[32] Mei Wang and Weihong Deng. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9322–9331, 2020.

[33] Mei Wang and Weihong Deng. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, 2021.

[34] Mei Wang, Weihong Deng, Jiani Hu, Xunqiang Tao, and Yaohai Huang. Racial faces in the wild: Reducing racial bias by information maximization adaptation network. In *ICCV*, pages 692–702. IEEE, 2019.

[35] Mei Wang, Yaobin Zhang, and Weihong Deng. Meta balanced network for fair face recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(11):8433–8448, 2022. doi: 10.1109/TPAMI.2021.3103191. URL <https://doi.org/10.1109/TPAMI.2021.3103191>.

[36] Haiyu Wu, Grace Bezold, Aman Bhatta, and Kevin W. Bowyer. Logical consistency and greater descriptive power for facial hair attribute learning. In *CVPR*, pages 8588–8597. IEEE, 2023.

[37] Haiyu Wu, Sicong Tian, Aman Bhatta, Kagan Öztürk, Karl Ricanek, and Kevin W. Bowyer. Facial hair area in face recognition across demographics: Small size, big effect. In *WACV (Workshops)*, pages 1131–1140. IEEE, 2024.

[38] Xingkun Xu, Yuge Huang, Pengcheng Shen, Shaoxin Li, Jilin Li, Feiyue Huang, Yong Li, and Zhen Cui. Consistent instance false positive improves fairness in face recognition. In *CVPR*, pages 578–586. Computer Vision Foundation / IEEE, 2021.

- [39] Michael Yeung, Toya Teramoto, Songtao Wu, Tatsuo Fujiwara, Kenji Suzuki, and Tamaki Kojima. Variface: Fair and diverse synthetic dataset generation for face recognition. *CoRR*, abs/2412.06235, 2024.
- [40] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, pages 11941–11952. IEEE, 2023.