

DeAR: Debiasing Vision-Language Models with Additive Residuals

Ashish Seth*
IIT Madras, India

cs20s030@smail.iitm.ac.in

Mayur Hemani*
Adobe Inc.

mayur@adobe.com

Chirag Agarwal
Adobe Inc.

cagarwal@adobe.com

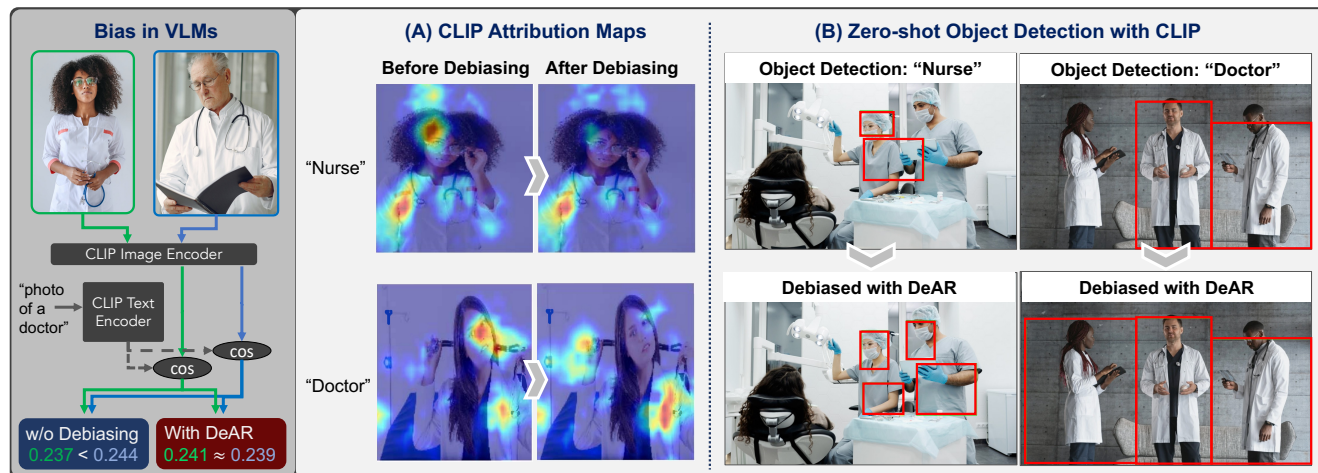


Figure 1. We present DEAR – a framework to de-bias large Vision-Language models (VLM) like CLIP [45], exhibited in the skewed similarity between specific language concepts and images of people of certain visual characteristics. (A) Attribution maps from the DEAR-augmented CLIP model indicate how the attribution for a text concept shifts from the person’s facial characteristics to objective cues in the image. (B) Results for zero-shot object detection with CLIP-ODS [48] that uses CLIP before and after debiasing show a clear improvement in the fairness of its detection results.

Abstract

Large pre-trained vision-language models (VLMs) reduce the time for developing predictive models for various vision-grounded language downstream tasks by providing rich, adaptable image and text representations. However, these models suffer from societal biases owing to the skewed distribution of various identity groups in the training data. These biases manifest as the skewed similarity between the representations for specific text concepts and images of people of different identity groups and, therefore, limit the usefulness of such models in real-world high-stakes applications. In this work, we present DEAR (Debiasing with Additive Residuals), a novel debiasing method that learns additive residual image representations to offset the original representations, ensuring fair output representations. In doing so, it reduces the ability of the representations to distinguish between the different identity groups. Further, we observe that the current fairness tests are performed

on limited face image datasets that fail to indicate why a specific text concept should/should not apply to them. To bridge this gap and better evaluate DEAR, we introduce the PROTECTED ATTRIBUTE TAG ASSOCIATION (PATA) dataset – a new context-based bias benchmarking dataset for evaluating the fairness of large pre-trained VLMs. Additionally, PATA provides visual context for a diverse human population in different scenarios with both positive and negative connotations. Experimental results for fairness and zero-shot performance preservation using multiple datasets demonstrate the efficacy of our framework. The dataset is released [here](#).

1. Introduction

Deep learning-based vision-language models (VLMs) [45] unify text and visual data into a common representation and reduce the computing cost of training for specific computer vision [19] and visually-grounded linguistic tasks [12, 47]. VLMs are trained with a large amount of data with the

*Equal Contribution

aim of matching image and text representations for image-caption pairs to capture diverse visual and linguistic concepts. However, VLMs exhibit societal biases manifesting as the skew in the similarity between their representation of certain textual concepts and kinds of images [2, 4, 31]. These biases arise from the underlying imbalances in training data [2, 3] and flawed training practices [55]. In this work, we present a method to significantly reduce the bias in VLMs by modifying the visual features of the models.

These societal biases in VLMs show up as the selective association (or dissociation) of their representations of human images with specific physical characteristics and their text representations for describing social labels [52]. For instance, a higher degree of similarity between the representation of the text “doctor” and images of men than that of the women can have trust consequences for models using the representations from these VLMs. To estimate the degree of such biases in VLMs, we compute the cosine similarity between the representations of a set of human images and specific key text phrases and compare their distribution over some associated *protected attributes*. These attributes represent visually discernible characteristics like *gender*, *race*, and *age* common to certain collective identity groups. In this work, we introduce DEAR, an additive residual-based de-biasing technique that can be augmented with any pre-trained visual-language model with separate visual and language encoders [33, 45, 49] to improve their fairness.

Our empirical analysis indicates that the protected attribute (PA) information from an image-text pair can be disentangled using their representation by simply learning a linear transformation of the visual representations produced by a VLM and subtracting (adding a negative residual) them from the original representation. We propose to train our framework using two objectives: i) learn a residual representation that when added to the original representation, renders it incapable of predicting different protected attributes, and ii) ensure that this modified representation is as close to the original as possible. We demonstrate that learning *additive residual* enables the de-biasing of pre-trained VLMs using quantitative skew computations on multiple datasets and qualitative evaluations. We also show that the resulting representation retains much of its predictive properties by means of zero-shot evaluation on different downstream tasks.

Recent efforts to mitigate these biases from VLMs, such as by Berg et al. [2], use unimodal datasets like the FairFace [28] and UTK [63] datasets. These face-image datasets lack the context necessary to infer the situations in which benign text associations can turn offensive when applied selectively. For instance, if the image of a person drinking water from a bottle is misconstrued as partaking of alcohol and if this association (in terms of image-text similarity) is selective to a specific group of people with

some specific visual characteristics, the association may be deemed offensive. To this end, we introduce the PROTECTED ATTRIBUTE TAG ASSOCIATION (PATA) dataset to test different associations for VLMs. PATA comprises images of persons in different contexts and their respective text captions with positive and negative connotations. We present our de-biasing results on both the FairFace and the PATA datasets. In summary, the paper makes the following contributions:

- We present the DEAR framework – a simple, computationally-efficient, and effective de-biasing method for VLMs that only adapts the image encoder of the VLM by adding a learned residual representation to it.
- We introduce the PROTECTED ATTRIBUTE TAG ASSOCIATION dataset – a novel context-based bias evaluation dataset that enables nuanced reporting of biases in VLMs for race, age, and gender protected attributes.

2. Related work

Our work lies at the intersection of large pre-trained VLMs and the identification, measurement, and mitigation of societal biases in large pre-trained models.

Pre-trained Vision-Language Models. Large VLMs aim to learn general information from a large amount of data and then transfer the knowledge to diverse downstream tasks [11, 29, 34, 37, 38, 45, 49, 56]. Recent works [14, 49, 58, 62] leverage contrastive strategies to learn representations for multimodal data. Based on how these models learn these representations, VLMs are broadly categorized into two groups: i) learning image-text representations jointly using transformer-based encoders [11, 14, 36, 37, 39, 49, 58, 62], and ii) learning individual unimodal encoders for image and text [5, 8–10, 15, 22, 46, 57]. Models like CLIP [45], ALIGN [23], and BASIC [43] are pre-trained on large datasets collected from the web to bring representations of paired images and text close to each other while distancing them from other pairs. Our work analyzes recently proposed state-of-the-art VLMs for the societal biases they exhibit. Specifically, we look at the behavior of CLIP [45], FLAVA [49], and BLIP [33] and propose to alleviate bias from their visual representations.

Fairness Techniques. Prior work in language [4, 20, 31], computer vision [17, 51], and graphs [1, 27, 40, 54] has primarily focused on debiasing models trained on unimodal data and are limited in scope as they only investigate gender bias, racial bias, or their intersections. In particular, their bias mitigation techniques can be broadly categorized into i) *pre-processing*, which modifies individual input features and labels [6], modifies the weights of the training samples [24], or obfuscates protected attribute information during the training process [59]; ii) *in-processing*, which uses adversarial techniques to maximize accuracy and re-

duce bias for a given protected attribute [61], data augmentation [1] or adding a bias-aware regularization term to the training objectives [26], and iii) *post-processing*, which changes the output predictions from predictive models to make them fairer [21, 25, 44]. It is non-trivial to apply these techniques to pre-trained VLMs because the training requires a large annotated dataset and immense computing resources. Mitigating bias in large pre-trained VLMs is a nascent research direction. Wang et al. [53] propose to remove the dimensions in CLIP embeddings most associated with gender bias, while Berg et al. [2] use adversarial fine-tuning with a corpus of face images and arrays of text prompts to mitigate bias. The former cannot achieve joint mitigation of bias for multiple protected attributes, and the latter method modifies the original model changing its accuracy on zero-shot tasks significantly. Our proposed DEAR framework addresses both issues by training a lightweight residual representation over the visual representations of VLMs, modeling the joint bias with respect to multiple protected attributes (gender, race, and age), and ensuring the similarity of the modified representation to the original.

Fairness Benchmarks. Previous work on debiasing VLMs [2, 53] exclusively focuses on face-image datasets, *i.e.*, FairFace [28] and UTKFace [63] datasets. While the FairFace dataset has 10,954 test images consisting of faces of 10 age groups, seven race groups, and the binary genders, the UTKFace dataset has 23,708 cropped test images in 104 different age values (1-105), five race groups, and the binary genders. Neither datasets have the pixels or annotations to provide the context of the person in an image. We collect a dataset of 4934 test images organized into different contextual scenes and prepare positive and negative captions for each context, providing a more nuanced view of a VLM’s fairness.

Next, we discuss how we achieve this accuracy-preserving joint mitigation of bias from multiple VLMs.

3. Preliminaries

We formally define the notion of fairness for VLMs, stipulate the fairness problem, and motivate the use of an additive residual for bias mitigation in VLMs.

Visually Discernible Protected Attributes (VDPAs). A *protected attribute* (PA) is defined as any characteristic of a group of people associated with their collective identity. A subset of these attributes are visually discernible, *i.e.*, it’s possible to label an image of a person belonging to a particular identity group with the corresponding protected label (PL). Our study includes the i) the binary *gender* of a person, ii) the ethno-racial characteristics (skin color, face shape, hair color, etc.) captured as *race*, and iii) the *age* of a person. For simplicity, we associate a single ground-truth VDPa label for an image \mathbf{I}_i , which has a gender label y_i^g , a race label y_i^r and an age label y_i^a . We denote a protected

attribute as $P(\mathbf{I})$ with the corresponding labels as a set $\{p_i\}$.

Abstract Model of Vision-Language Models. An abstract VLM consists of two key components: i) an image encoder E_i from the space of a 2D, 3-channel image to a latent space in \mathbb{R}^d and ii) a text encoder E_t from the space of text tokens to a similar latent space in \mathbb{R}^d . For an image \mathbf{I} and text token sequence \mathbf{T} , the output of the image and text encoders are denoted as $E_i(\mathbf{I})$ and $E_t(\mathbf{T})$. Finally, a fusion function combines the encoded representations from E_i and E_t into a common representation, which we disregard because of the manner in which they compute similarity.

Similarity of Image and Text Representations. Most large pre-trained VLMs [33, 46, 49] learn image and text representations that are in a similar latent space. We compute the similarity between an image (\mathbf{I}) and a text sequence (\mathbf{T}) using the cosine similarity between their representations:

$$\text{Sim}_E(\mathbf{I}, \mathbf{T}) = \frac{E_i(\mathbf{I}) \cdot E_t(\mathbf{T})^\top}{\|E_i(\mathbf{I})\| \|E_t(\mathbf{T})\|}, \quad (1)$$

Formally, the VLM encoders are fair for a protected attribute if the similarity between an image \mathbf{I} and a text sequence \mathbf{T} is independent of the protected label for \mathbf{I} . It is formalized by measuring the skew of the similarity scores of \mathbf{T} with a set of images, across their protected labels.¹

Fairness and Fair Encoders. We denote the probability of randomly selecting an image \mathbf{I} from a set of images \mathcal{I} such that it has a protected attribute label $P(\mathbf{I}) = p_i$ with f_i . Next, we define a matching function $M_\epsilon(\mathbf{I}, \mathbf{T})$ which is 1 if $\text{Sim}_E(\mathbf{I}, \mathbf{T}) \geq \epsilon$, and 0 otherwise. We define \mathcal{I}^m as the subset of \mathcal{I} for which $M_\epsilon(\mathbf{I}, \mathbf{T})$ is 1 for each $\mathbf{I} \in \mathcal{I}^m$ and the fraction of these images that also have the protected label p_i is denoted by f_i^m . The ratio of f_i^m and f_i is considered as the skew in the distribution of images with the protected label p_i imposed by the model (E_i, E_t) by virtue of the similarity between its output representations. The *Skew* measure for the i^{th} protected label, with respect to the text sequence and image set, is then computed as:

$$\text{Skew}_i(\mathcal{I}, \mathbf{T}) = \log(f_i^m / f_i), \quad (2)$$

where the score is zero when $f_i^m = f_i$, *i.e.*, the matching does not depend on the value of the protected attribute.

The goal of learning a fair encoding then consists of training encoders (E_i^*, E_t^*) such that $\text{Skew}_i(\mathcal{I}, \mathbf{T}) \approx 0$ for all text sequences \mathbf{T} , protected attributes i , and for any set of images \mathcal{I} . We posit that a model devoid of any information useful for distinguishing between different PLs for a PA is fair because encoding from such a model is independent of that particular PA. We submit that learning such a model with respect to a VLM is possible by adapting the visual encoder of the VLM. This is because the encoding from it

¹The formulation used here is adapted from [18].

captures rich information in a compact form, and it is easier to find training data to adapt the image encoder than it is for the text encoder. We posit that adapting just the image encoder can produce a fair model. We corroborate this position next with empirical evidence.

Disentangling Protected Attribute Information. To motivate the design of DEAR, we present empirical evidence for the bias problem using the `gender` attribute. We show that $E_i(\mathbf{I})$ can be split into two additive components - i) $\phi(\mathbf{I})$ that represents the gender information for \mathbf{I} and ii) a protected-attribute (PA) free representation $\bar{\phi}(\mathbf{I})$ that cannot be used to identify the gender of the person in the image:

$$E_i(\mathbf{I}) = \bar{\phi}(\mathbf{I}) + \phi(\mathbf{I}) \quad (3)$$

For the protected labels `male` (m) and `female` (f), we get:

$$E_i(\mathbf{I}_f) = \bar{\phi}(\mathbf{I}_f) + \phi(\mathbf{I}_f) \quad (4)$$

$$E_i(\mathbf{I}_m) = \bar{\phi}(\mathbf{I}_m) + \phi(\mathbf{I}_m) \quad (5)$$

$$\Rightarrow E_i(\mathbf{I}_f) - E_i(\mathbf{I}_m) = \phi(\mathbf{I}_f) - \phi(\mathbf{I}_m), \quad (6)$$

where $\bar{\phi}(\mathbf{I}_f) = \bar{\phi}(\mathbf{I}_m)$ as both contain the PA free representation based on the context in the image. Further, we hypothesize that $\phi(\mathbf{I})$ can be adequately captured for a VLM by the corresponding text encoding for the equivalent textual description “photo of a man/woman” (t_m/t_f) using a linear transformation. This is expressed as:

$$\begin{aligned} \phi(\mathbf{I}_f) &= \mathbf{A} \cdot E_t(t_f) + \mathbf{B} \\ \phi(\mathbf{I}_m) &= \mathbf{A} \cdot E_t(t_m) + \mathbf{B}, \end{aligned} \quad (7)$$

where $\mathbf{A} \in \mathbb{R}^{d \times d}$ is a transformation matrix and $\mathbf{B} \in \mathbb{R}^d$ is a vector of the same dimensions as the representations. Combining Eqns. 6 and 7, we get:

$$\begin{aligned} E_i(\mathbf{I}_f) - E_i(\mathbf{I}_m) &= \mathbf{A} \cdot E_t(t_f) - \mathbf{A} \cdot E_t(t_m) \\ \Rightarrow \mathbf{K}[E_i(\mathbf{I}_f) - E_i(\mathbf{I}_m)] &= E_t(t_f) - E_t(t_m), \end{aligned} \quad (8)$$

where $\mathbf{K} = \mathbf{A}^{-1}$. If we can find evidence that such a linear transformation \mathbf{K} of the difference of image representation exists that matches the difference between the encoding of the corresponding text concepts, then we can learn a linear function that disentangles $\phi(\mathbf{I})$ from $E_i(\mathbf{I})$.

Empirical evidence: We train a regression model on 10000 random pairs of images of women and men in a professional workplace setting (a subset of the FairFace [28] dataset), to learn a linear function from the difference of the CLIP image representations to the difference of the CLIP text representations for the “photo of a woman” and “photo of a man” phrases. We found the model to have a bounded relative mean-squared error ($<10^{-6}$) over 100 repetitions of the experiment for unseen pairs, which demonstrates that a piece-wise linear function can possibly be trained to approximate $\phi(\mathbf{I})$. Next, we describe how our DEAR framework trains to approximate $\phi(\mathbf{I})$ for multiple protected attributes, and uses it to remove protected attribute information from $E_i(\mathbf{I})$.

4. Our Framework: DEAR

Our framework comprises of two key components: i) an Additive Residual Learner (ARL) component that learns to disentangle the protected attribute information $\phi(\mathbf{I})$ from the image representation produced by the visual encoder of a given VLM and ii) the protected attribute misclassification objective designed to train the ARL component. We first discuss the ARL component and then present details of the training network and objectives.

4.1. Additive Residual Learning

As discussed in Sec. 3, we propose to train a network $R(\cdot)$ that outputs a representation $\phi(\mathbf{I})$ so that we can compute the protected-attribute free representation $\bar{\phi}(\mathbf{I})$:

$$\bar{\phi}(\mathbf{I}) = E_i(\mathbf{I}) - R(E_i(\mathbf{I})) \quad (9)$$

For brevity, we denote $R(E_i(\mathbf{I}))$ as $R(\mathbf{I})$. Note that the network $R: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a single linear transformation, where d is the number of dimensions of the image representation produced by E_i . This is followed by a function $\sigma_{\text{act}}(\cdot)$ that represents the combination of activation and regularization layers (dropout, etc.) as used by the last layer of the image encoder E_i . If the last layer does not use any activation function, σ_{act} is an identity transformation, *i.e.*,

$$\bar{\phi}(\mathbf{I}) = E_i(\mathbf{I}) + \sigma_{\text{act}}(R(\mathbf{I})), \quad (10)$$

where $\bar{\phi}(\mathbf{I})$ represents the protected-attribute free representation for image \mathbf{I} . Note that the reason for transforming the last layer activation and regularization mechanisms from the image encoder is to align the embedding space, where $\bar{\phi}(\mathbf{I})$ maps to the original encoding $E_i(\mathbf{I})$ so that the image-text equivalence still holds. The resulting representation has two primary goals: i) it matches the original representation so that the zero-shot accuracy of the model can be retained, *i.e.*, $\bar{\phi}(\mathbf{I}) \approx E_i(\mathbf{I})$, and ii) it cannot be used to distinguish the protected label for the image with respect to different protected attribute. The above goals are achieved through different training objectives, which we describe next.

4.2. Training ARL with Protected Attribute Misclassification

To train the ARL to produce representations that cannot be used to identify the protected labels for an image, we pre-train an adversary – the Protected Attribute Classifier (PAC) network and use its classification loss as an objective to be maximized for training the ARL. Unlike previous works in bias mitigation that jointly trains an adversarial network, we separately train the PAC. This is possible because the PAC can be trained with the underlying VLM representations and be used to supply gradients as the representations change to their protected-attribute-free form during training.

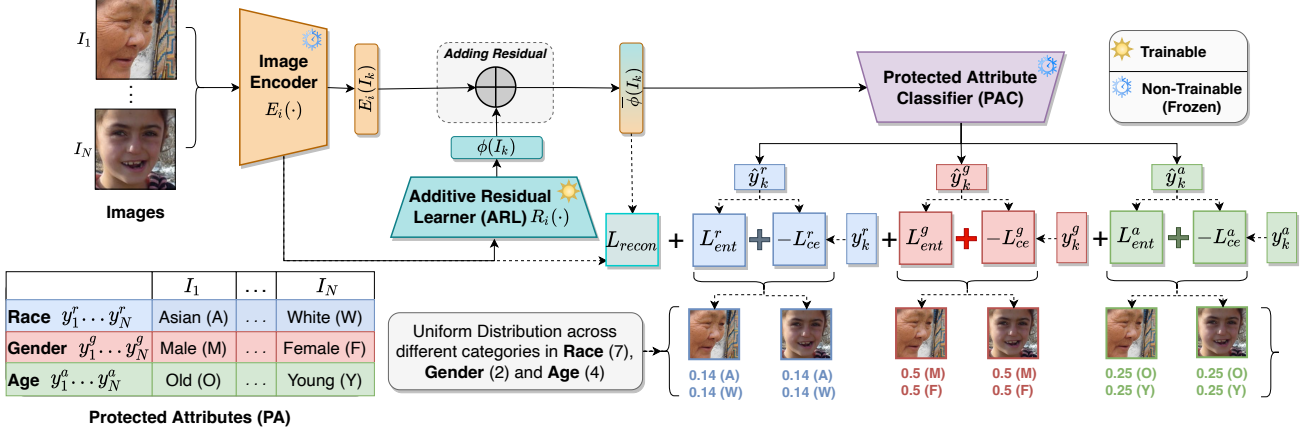


Figure 2. Our DEAR framework can learn fair representations for pre-trained Vision-Language models. DEAR is a debiasing framework that takes image representations $E_i(I_k)$ as input and returns a residual representation $\phi(I_k)$. Next, it adds $\phi(I_k)$ with $E_i(I_k)$ to get debiased representation $\bar{\phi}(I_k)$ that is fed to a protected attribute classifier to get prediction for multiple protected attributes (race \hat{y}_k^r , gender \hat{y}_k^g , and age \hat{y}_k^a). DEAR jointly minimizes the negative of Cross Entropy loss for race (L_{ce}^r), gender (L_{ce}^g), and age (L_{ce}^a) attributes for debiasing across multiple protected attributes and a reconstruction loss L_{recon} to retain the zero-shot performance of the underlying VLM.

The PAC module uses ℓ_2 normalized d -dimensional representation from the VLM image encoder E_i as input and generates multiple sets of logits for individual protected attributes. It consists of a single linear layer ($d \times 256$) with ReLU activation, followed by multiple linear classification projection heads (256×128) with ReLU, and a linear layer to produce logits with output sizes 7 (for race), 4 (for age) and 2 (for gender), where the logit sizes are determined by the FairFace [28] training dataset. A unified architecture like DEAR helps in jointly modeling multiple protected attributes associated with an image. The network is trained on the combined cross-entropy loss using the ground-truth protected labels (race, age, and gender) for each image:

$$\mathcal{L}_{PAC} = \mathcal{L}_{ce}^r(y^r, \hat{y}^r) + \mathcal{L}_{ce}^g(y^g, \hat{y}^g) + \mathcal{L}_{ce}^a(y^a, \hat{y}^a), \quad (11)$$

where \hat{y}^r , \hat{y}^g and \hat{y}^a are the predictions, and y^r , y^g and y^a are the ground-truth labels for the protected attributes race, gender, and age respectively.

Training objectives for the ARL module. The VLM parameters are frozen while the ARL module is trained using three objectives: i) minimizing the reconstruction loss L_{recon} , which is the mean ℓ_2 loss between the unmodified image representations $E_i(\mathbf{I})$ and the modified protected attribute information-free $\bar{\phi}(\mathbf{I})$, ii) the $L_{ent}^{\{r,g,a\}}$ to minimize the maximum soft-max probability for each protected attribute classifier head, and iii) $L_{ce}^{r,g,a}$, to maximize the misclassification of the protected label by the PAC module.

$$L_{ARL} = w_{recon} \cdot L_{recon} + w_{ent} \cdot (L_{ent}^r + L_{ent}^g + L_{ent}^a) - (w_{ce}^r \cdot L_{ce}^r + w_{ce}^a \cdot L_{ce}^a + w_{ce}^g \cdot L_{ce}^g), \quad (12)$$

where $L_{recon} = \|E_i(\mathbf{I}) - \bar{\phi}(\mathbf{I})\|_2$, $L_{ent}^x = \max(\sigma_{softmax}(\hat{y}_x))$ (for PA x), $\sigma_{softmax}(x) = \text{Softmax}(x)$, and the w_i 's are

the regularization weights that assist the training of DEAR across multiple protected attributes.

The negative cross-entropy losses aim to maximize the cross-entropy loss for the PAC. However, when used alone, this loss results in the protected label predictions getting flipped (from one label to another) rather than reducing the decision confidence of the PAC. The L_{ent} ensures that the PAC module produces high-entropy logits from the modified image representation $\bar{\phi}(\mathbf{I})$ so that it cannot be easily used for predicting the protected attributes with high confidence. Figure 2 describes the operation of the DEAR framework in detail. The success of the network is evaluated by the degree to which the i) bias in the model is mitigated and ii) zero-shot performance of the model is retained.

5. Experiments

Here, we describe the datasets, the metrics, and the implementation parameters for our experiments.

Datasets. We use two datasets each consisting of a set of images and their ground truth (protected labels) to analyze the bias in the output representations of VLMs. Further, each image in the datasets is tested against a set of curated text captions that represent different desirable (positive) and undesirable (negative) connotations.

Images and Protected Labels. We obtain images and labels from the following: i) the *FairFace dataset* [28] consists of a set of cropped face images of different persons. It has 86,700 training images and 10,954 validation images, annotated with the binary gender (male and female), ethnic-racial classes (White, Black, Indian, East Asian, Southeast Asian, Middle Eastern, and Latino), and multiple age brackets, which we bin into four categories (*child* < 20, *young*

< 40, *middle-aged* < 60, and *senior* > 60), and ii) the PROTECTED ATTRIBUTE TAG ASSOCIATION (PATA) dataset consists of 4,934 public images crawled from the web, organized in 24 scenes each with between 100 and 400 images. These scenes represent situations in which humans of specific identity groups may have biased notions associated with them. The images are further annotated with the binary gender, ethnic-racial labels (White, Black, Indian, East Asian, and Latino-Hispanic), and two age group labels (young and old adults).

Text Captions. Each scene in the PATA dataset is associated with the positive caption set consisting of desirable descriptions of a person, designed specifically to elicit a biased response from the models. For instance, for the `race` protected attribute (PA), the caption “photo of a law-abiding citizen”, and for the `gender` protected attribute, the caption “photo of a smart engineer” must both result in uniformly distributed similarities between the images and the text. The negative captions similarly consist of untoward and offensive text in the context of each protected attribute. For the PATA dataset, there are 78 (unique) positive captions, and 88 negative captions, with an average of 1.08 and 1.2 per (scene, PA) pairs, respectively. For the FairFace dataset, we consider all images to belong to a single scene and use 16 positive and negative captions for each.

Advantages of the PATA dataset. Images in the FairFace dataset only have cropped faces and, therefore, do not contain any context. Hence, any association between the text captions and the face images is incidental at best because the neural network does not have any signal to form the association. The PATA dataset bridges this gap and thus enables a more nuanced evaluation of the bias presented by a VLM. More details about the dataset are included in the Appendix A of the supplemental material.

Zero-shot Evaluation datasets. We consider the CIFAR-10, CIFAR-100 [32], FER2013 [60], and ImageNet [13] datasets for zero-shot image classification task using debiased VLM representations. Further, we consider the UCF-101 [50], Kinetics-700 [7], and RareAct [41] datasets for zero-shot video classification.

Evaluation Metrics. Our de-biasing framework is evaluated on two key criteria: i) the degree to which it improves the fairness of a given VLM, which is measured using the similarity between person images with different protected labels and a set of text captions, and ii) the degree to which it retains the zero-shot performance of the underlying VLM on various tasks, including classification on human and non-human related datasets.

Fairness Metrics. The degree of fairness exhibited by a model is measured in two respects:

i) the overall shift in similarity characteristics of the image and text features towards fairness measured using the mean *MaxSkew* and the mean *MinSkew* scores. The former indi-

cates the selective association of the model in favor of any protected label and the latter indicates the degree of selective dissociation against any other protected label. These are computed for a set of images \mathbf{I} , a set of captions \mathbf{T} , and a protected attribute P with labels p_i as:

$$\begin{aligned}\psi_{max}(\mathbf{I}, \mathbf{T}, P) &= \text{Mean}_{t \in \mathbf{T}} [\max_i (\text{Skew}_i(\mathbf{I}, t))] \\ \psi_{min}(\mathbf{I}, \mathbf{T}, P) &= -\text{Mean}_{t \in \mathbf{T}} [\min_i (\text{Skew}_i(\mathbf{I}, t))],\end{aligned}\tag{13}$$

where $\text{Skew}_i(\cdot)$ is computed using Equation 2.

ii) the ordered measurement of the similarity distribution across the protected labels of each protected attribute, which is indicated using the *MaxSkew@K* and *MinSkew@K* scores as stipulated in Geyik et al. [18] and used for fairness analysis. This is a ranking-based measurement that only considers the k most similar images to each text caption, and then computes the mean Max- and Min-Skew scores. To add more nuance to the evaluation, we compute the above metrics for both positive and negative captions for the person in the image, which clearly indicates the stakes for the model in selective association or dissociation for a particular identity group.

Zero-shot performance metrics. We evaluate the zero-shot performance of DEAR-augmented and vanilla VLMs using standard classification accuracy metrics for the image-classification tasks. For video action recognition tasks, we use the top-1 accuracy, the averaged accuracy measures and variants of the mean weighted average precision metrics as described in Miech et al. [42]. We report the drop in each metric to indicate how much the debiasing process affects the zero-shot performance of the underlying VLM. In addition, we discuss the effect of DEAR framework on the task of zero-shot open-vocabulary object detection [48].

Implementation Details. Both PAC and the ARL modules are individually trained for each VLM using the training subset of the FairFace dataset, which consists of 86,700 images with annotations for the race, gender, and age attributes. The PAC is trained on its objectives (Equation 11) for 10 epochs, whereas the ARL is trained for 30 epochs with early stopping based on validation loss saturation. See Appendix B.3 for more details on other hyperparameters.

6. Results

Next, we present the experimental results for the DEAR framework and address the following key questions: Q1) Does DEAR reduce the bias in VLMs? and Q2) Do the de-biased VLMs retain their predictive properties?

Q1) Bias Reduction in VLMs. We evaluate both the overall shift in the distribution of the similarities (*MaxSkew*) and the ordered metrics (*MaxSkew@K*) between the images and text captions for both evaluation datasets. Tables 1-2 present the mean Max-/MinSkew and mean Max-/MinSkew@K scores for the FairFace and PATA datasets.

Table 1. Systematic bias evaluation of VLMs on the **Fair-Face** dataset and their DEAR counterparts using *MaxSkew*, *MinSkew*, $MS@k=MaxSkew@k$, $mS@k=MinSkew@k$ metrics on FairFace dataset. {+/-} refers to the positive and negative sentiments. [C]=CLIP, [C]_D=DEAR-CLIP, [F]=FLAVA, [F]_D=DEAR-FLAVA. Values closer to zero indicate better fairness. DEAR-augmented VLMs exhibit better fairness. See Appendix C for complete results on other VLMs.

PA	+/-	MaxSkew		MinSkew		MaxSkew		MinSkew	
		[C]	[C] _D	[C]	[C] _D	[F]	[F] _D	[F]	[F] _D
Age	+ve	0.36	0.10	-0.25	-0.09	0.33	0.08	-0.34	-0.35
	-ve	0.54	0.45	-0.53	-0.43	0.48	0.20	-0.37	-0.26
Race	+ve	0.49	0.33	-0.85	-0.69	0.31	0.37	-0.54	-0.45
	-ve	0.59	0.55	-2.49	-0.87	0.38	0.26	-0.54	-0.42
Gender	+ve	0.29	0.19	-0.60	-0.32	0.25	0.14	-0.44	-0.22
	-ve	0.22	0.21	-0.41	-0.34	0.29	0.14	-0.57	-0.18
		MS@k		mS@k		MS@k		mS@k	
Age	+ve	0.48	0.14	-0.22	-0.08	0.42	0.23	-3.47	-3.41
	-ve	0.71	0.38	-2.58	-2.50	0.51	0.31	-2.84	-2.36
Race	+ve	0.55	0.40	-3.19	-0.61	0.48	0.59	-2.70	-2.89
	-ve	0.84	0.62	-5.12	-6.28	0.45	0.38	-0.79	-1.85
Gender	+ve	0.29	0.17	-0.57	-0.29	0.21	0.12	-2.38	-0.14
	-ve	0.36	0.14	-7.71	-0.22	0.30	0.17	-2.90	-2.62

Both are computed using a threshold cosine similarity value of 0.1 for both datasets. We chose $k=1000$ for FairFace and $k=100$ for the PATA dataset.

Interpretation. The mean MaxSkew and MaxSkew@K scores indicate the selective association with respect to some favored protected labels for each protected attribute. The mean MinSkew and MinSkew@K scores indicate selective dissociation of the models for particular protected labels. Note that the favored (or disfavored) labels can change between the original and de-biased variants of each model. If the sentiment of an entry with a high MaxSkew value is negative, the favored label is the one with a disproportionately high association with negative captions. Similarly, entries with a low (negative) MinSkew value for a positive sentiment entry imply that the disfavored label has an unusually low association with the positive captions in the dataset.

Both tables indicate the benefit of using the DEAR framework for de-biasing. We observe a slight increase in the skew values in some cases which we attribute to a flip in the favored PA labels. For instance, the MaxSkew value for FLAVA changes from 0.31 to 0.37 for the (race, positive) combination, but so does the favored label from “middle-eastern” to “east-asian”, and going overboard with the optimization. This can be mitigated by reducing the associated training weights w_{ce}^r while training the ARL module.

Q2) Zero-Shot Accuracy Evaluation. VLMs de-biased with DEAR retain much of their zero-shot performance on downstream tasks. We perform an extensive analysis of the effect of our DEAR framework on several downstream

Table 2. Bias evaluation of VLMs and their DEAR counterparts on the **PATA** dataset using *MaxSkew*, *MinSkew*, $MS@k=MaxSkew@k$, $mS@k=MinSkew@k$ metrics on FairFace dataset. {+/-} refers to the positive and negative sentiments. [C]=CLIP, [C]_D=DEAR-CLIP, [F]=FLAVA, [F]_D=DEAR-FLAVA. Values closer to zero indicate better fairness. DEAR-augmented VLMs exhibit better fairness. See Appendix C for complete results on other VLMs.

PA	+/-	MaxSkew		MinSkew		MaxSkew		MinSkew	
		[C]	[C] _D	[C]	[C] _D	[F]	[F] _D	[F]	[F] _D
Age	+ve	0.10	0.07	-0.12	-0.08	0.06	0.05	-0.08	-0.05
	-ve	0.19	0.12	-1.21	-0.16	0.22	0.15	-0.45	-0.23
Race	+ve	0.15	0.15	-0.41	-0.40	0.11	0.12	-0.14	-0.14
	-ve	0.43	0.38	-3.20	-2.14	0.34	0.33	-2.46	-2.49
Gender	+ve	0.085	0.07	-0.11	-0.09	0.07	0.06	-0.08	-0.08
	-ve	0.20	0.18	-0.76	-0.70	0.16	0.14	-0.61	-0.22
		MS@k		mS@k		MS@k		mS@k	
Age	+ve	0.10	0.08	-0.13	-0.10	0.10	0.07	-0.13	-0.08
	-ve	0.16	0.12	-0.22	-0.14	0.25	0.23	-1.84	-4.50
Race	+ve	0.25	0.22	-0.63	-0.32	0.23	0.26	-0.29	-0.73
	-ve	0.66	0.61	-5.70	-4.12	0.48	0.51	-3.20	-3.40
Gender	+ve	0.13	0.10	-0.16	-0.13	0.12	0.09	-0.22	-0.11
	-ve	0.24	0.22	-2.59	-2.13	0.20	0.15	-1.40	-0.58

Table 3. Results of state-of-the-art visual-language models and their DEAR counterparts for four image classification datasets. Across five pre-trained visual-language models, DEAR achieves zero-shot performance similar to vanilla models. See Appendix C for complete results.

Model	C-10	C-100	FER2013	ImageNet
CLIP (ViT/B-32)	89.93	62.93	43.83	58.08
DEAR-CLIP (ViT/B-32)	88.85	60.08	39.60	55.84
FLAVA	90.53	65.60	28.36	49.30
DEAR-FLAVA	89.05	64.00	27.19	47.67
BLIP	85.00	51.61	39.50	32.57
DEAR-BLIP	81.20	48.90	36.50	29.94

zero-shot classification tasks. In Tables 3-4, we present the zero-shot performance for four image classification and three video classification datasets. For general object classification datasets like CIFAR-10, CIFAR-100, and ImageNet, we observe that the zero-shot performance of DEAR augmented VLMs, on average, is only 2.2% lower than their vanilla counterparts. On a fine-tuned facial emotion classification dataset like FER2013, the zero-shot performance between vanilla VLMs and their DEAR counterpart is smaller (1.7%). Following [45], we sub-sample each video to only a single center frame for video classification, effectively turning it into an image classification dataset. On UCF-101 and Kinetics-700 datasets, we observe a drop of 2.06% and 1.77% for the DEAR VLMs in the zero-shot predictive performance. Finally, for the RareAct dataset that is used to

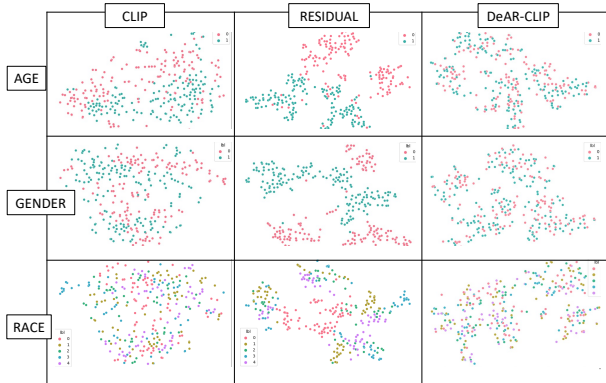


Figure 3. TSNE Plots for CLIP, Residual and DEAR-CLIP features for a subset of the PATA dataset indicate that the residual plots indeed capture the specific attributes and that the DEAR-CLIP features have greater overlap between points of different protected labels than the original features. Refer to Appendix C for higher resolution plot.

Table 4. Results of state-of-the-art visual-language models and their DEAR counterparts for three video classification datasets. Across five pre-trained visual-language models, DEAR achieves zero-shot performance similar to vanilla models. See Appendix C for complete results.

Model	UCF-101 Kinetics-700		RareAct	
	Top-1	AVG	mWAP	mWSAP
CLIP (Base-32)	57.65	43.97	16.63	16.78
DEAR-CLIP (Base-32)	55.77	42.20	16.02	16.03
FLAVA	39.09	37.85	16.12	16.14
DEAR-FLAVA	37.27	35.59	15.30	15.43
BLIP	43.26	37.07	16.35	16.44
DEAR-BLIP	40.34	34.78	15.86	15.92

quantify the zero-shot recognition of unusual actions like “blend phone” and “microwave shoes”, we observe a drop of 0.71% and 0.69% in the zero-shot classification performance of DEAR using mWAP and mWSAP metrics. Please refer to Appendix C for more details.

7. Ablation Studies and Analysis

We explore two important issues here – i) ascertaining that the de-biased representations are free from the protected attribute information and ii) justifying the joint training of the ARL module for different protected attributes. Please refer to Appendix D for more ablation studies.

Evidence for de-biasing. We run several experiments to ascertain that the representations from the DEAR-CLIP model are indeed free from the protected attribute information. Figure 1 (A) shows a comparison of the explanations for CLIP and its de-biased variant (DEAR-CLIP). These explanations are generated for some in-the-wild creative-

commons stock images using the approach in [16]. They illustrate how the original CLIP model focuses on the face and hair region for certain occupation-related keywords, and after de-biasing, the focus shifts to more indicative cues like the stethoscope for a doctor. Figure 1(B) shows results of zero-shot open-vocabulary object detection using [48]. Again, the base CLIP model exhibits a bias omitting the woman for the “doctor” keyword, and only detecting the female medical professional as a “nurse”. With DEAR-CLIP, both men and women are detected for both keywords. It is noteworthy that at a lower detection threshold, both male and female doctors and nurses may have been detected, but it still points to the implicit bias because the stereotypical associations rank higher in similarity. Figure 3 depicts how the CLIP features get disentangled into the target PAs, and the addition of the residual results in coinciding points for different protected labels, indicating that the representations for different PLs cannot be distinguished from each other.

Joint vs. Sequential training of ARL. Another design choice is to use a jointly-trained PAC module for training the ARL module. It is also possible to train the residual module using multiple individual classifiers (one for gender, one for race, etc.) and use different weights for the ARL modules for each protected attribute. At inference, the VLM’s image representation is transformed using multiple ARL weights one after the other. We experiment with this strategy for CLIP and find that while this strategy does reasonably well for debiasing (DEAR-CLIP: MaxSkew (‘gender’=0.75, ‘race’=0.167); Sequential strategy: MaxSkew (‘gender’=0.87, ‘race’=0.161), its zero-shot ImageNet accuracy considerably falls (DEAR-CLIP=55.84%, Sequential strategy=44.43%). Hence, the joint training approach is better than sequential strategy.

8. Conclusion

We address the problem of mitigating biases from the representations of large pre-trained visual-language models. To this end, we present DEAR (Debiasing with Additive Residuals), a novel method that learns additive residual image representations to offset the original VLM representations, ensuring fair output representations. In particular, DEAR comprises an additive residual learner component, which learns to disentangle the protected attribute information from the output representations. We also introduce a new context-based bias benchmarking dataset for VLMs - the Protected Attribute Tag Association (PATA) dataset. Our results on multiple benchmarking datasets show that DEAR can significantly improve the fairness of output representations without sacrificing predictive zero-shot performance. With the increasing development of large VLM, our work paves the way for several exciting future directions for fairness research.

References

- [1] Chirag Agarwal, Himabindu Lakkaraju, and Marinka Zitnik. Towards a unified framework for fair and stable graph representation learning. In *UAI*, 2021. 2, 3
- [2] Hugo Berg, Siobhan Mackenzie Hall, Yash Bhalgat, Wonsuk Yang, Hannah Rose Kirk, Aleksandar Shtedritski, and Max Bain. A prompt array keeps the bias away: Debiasing vision-language models with adversarial learning. In *AAACL*, 2022. 2, 3, 12, 13
- [3] Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *CoRR*, abs/2110.01963, 2021. 2
- [4] Conrad Borchers, Dalia Sara Gala, Benjamin Gilbert, Eduard Oravkin, Wilfried Bounsi, Yuki M Asano, and Hannah Rose Kirk. Looking for a handsome carpenter! debiasing gpt-3 job advertisements. *ACL Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, 2022. 2
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, 2020. 2
- [6] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. Optimized pre-processing for discrimination prevention. In *NeurIPS*, 2017. 2
- [7] Joao Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *arXiv*, 2019. 6
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020. 2
- [9] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv*, 2020. 2
- [10] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, 2021. 2
- [11] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *ECCV*, 2020. 2
- [12] Jaemin Cho, Seunghyun Yoon, Ajinkya Kale, Franck Dernoncourt, Trung Bui, and Mohit Bansal. Fine-grained image captioning with CLIP reward. In *NAACL 2022*. 1
- [13] 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, 2009. 6
- [14] Karan Desai and Justin Johnson. Virtex: Learning visual representations from textual annotations. In *CVPR*, 2021. 2
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv*, 2018. 2
- [16] Ruth C Fong and Andrea Vedaldi. Interpretable explanations of black boxes by meaningful perturbation. In *ICCV*, 2017. 8
- [17] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 2021. 2
- [18] Sahin Cem Geyik, Stuart Ambler, and Krishnamurthy Kenthapadi. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. *CoRR*, abs/1905.01989, 2019. 3, 6
- [19] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *ICLR*, 2022. 1
- [20] Wei Guo and Aylin Caliskan. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. In *AIES*, 2021. 2
- [21] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *NeurIPS*, 2016. 3
- [22] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 2
- [23] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 2
- [24] Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems*, 2012. 2
- [25] Faisal Kamiran, Asim Karim, and Xiangliang Zhang. Decision theory for discrimination-aware classification. In *ICDM*, 2012. 3
- [26] Toshihiro Kamishima, Shotaro Akaho, Hideki Asoh, and Jun Sakuma. Fairness-aware classifier with prejudice remover regularizer. In *Joint European conference on machine learning and knowledge discovery in databases*, 2012. 3
- [27] Jian Kang and Hanghang Tong. Fair graph mining. In *ACM International Conference on Information & Knowledge Management*, 2021. 2
- [28] Kimmo Karkkainen and Jungseock Joo. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *WACV*, 2021. 2, 3, 4, 5, 11, 12
- [29] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. 2
- [30] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. 11
- [31] Hannah Rose Kirk, Yennie Jun, Filippo Volpin, Haider Iqbal, Elias Benussi, Frederic Dreyer, Aleksandar Shtedritski, and Yuki Asano. Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models. *NeurIPS*, 2021. 2
- [32] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [33] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 2, 3, 12

- [34] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 2021. 2
- [35] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *NeurIPS*, 2021. 12, 15
- [36] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv*, 2019. 2
- [37] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020. 2
- [38] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In *ICLR*, 2022. 2
- [39] Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 12-in-1: Multi-task vision and language representation learning. In *CVPR*, 2020. 2
- [40] Jing Ma, Ruocheng Guo, Mengting Wan, Longqi Yang, Aidong Zhang, and Jundong Li. Learning fair node representations with graph counterfactual fairness. In *ACM International Conference on Web Search and Data Mining*, 2022. 2
- [41] Antoine Miech, Jean-Baptiste Alayrac, Ivan Laptev, Josef Sivic, and Andrew Zisserman. Rareact: A video dataset of unusual interactions. *arXiv preprint arXiv:2008.01018*, 2020. 6
- [42] Antoine Miech, Jean-Baptiste Alayrac, Ivan Laptev, Josef Sivic, and Andrew Zisserman. Rareact: A video dataset of unusual interactions. *arxiv:2008.01018*, 2020. 6
- [43] Hieu Pham, Zihang Dai, Golnaz Ghiasi, Hanxiao Liu, Adams Wei Yu, Minh-Thang Luong, Mingxing Tan, and Quoc V Le. Combined scaling for zero-shot transfer learning. *arXiv*, 2021. 2
- [44] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. In *NeurIPS*, 2017. 3
- [45] 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 *ICML*, 2021. 1, 2, 7, 11, 12
- [46] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 2, 3
- [47] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can clip benefit vision-and-language tasks? *arXiv preprint arXiv:2107.06383*, 2021. 1
- [48] Alex Shonenkov, Sergey Shtekhin, and Denis Karachev. Clip object detection & segmentation. <https://github.com/shonenkov/CLIP-ODS>. 1, 6, 8
- [49] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. In *CVPR*, 2022. 2, 3
- [50] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 6
- [51] Angelina Wang, Alexander Liu, Ryan Zhang, Anat Kleiman, Leslie Kim, Dora Zhao, Iroha Shirai, Arvind Narayanan, and Olga Russakovsky. Revise: A tool for measuring and mitigating bias in visual datasets. *IJCV*, 2022. 2
- [52] Jialu Wang, Yang Liu, and Xin Eric Wang. Are gender-neutral queries really gender-neutral? mitigating gender bias in image search. In *EMNLP*, 2021. 2
- [53] Jialu Wang, Yang Liu, and Xin Eric Wang. Are gender-neutral queries really gender-neutral? mitigating gender bias in image search. In *EMNLP*, 2021. 3
- [54] Nan Wang, Lu Lin, Jundong Li, and Hongning Wang. Unbiased graph embedding with biased graph observations. In *WWW*, 2022. 2
- [55] Tianlu Wang, Jieyu Zhao, Mark Yatskar, Kai-Wei Chang, and Vicente Ordonez. Balanced datasets are not enough: Estimating and mitigating gender bias in deep image representations. In *International Conference on Computer Vision (ICCV)*, October 2019. 2
- [56] Wenhui Wang, Hangbo Bao, Li Dong, and Furu Wei. VLMo: Unified vision-language pre-training with mixture-of-modality-experts. 2021. 2
- [57] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *ICML*, 2021. 2
- [58] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. Merlot: Multimodal neural script knowledge models. *NeurIPS*, 2021. 2
- [59] Rich Zemel, Yu Wu, Kevin Swersky, Toni Pitassi, and Cynthia Dwork. Learning fair representations. In *ICML*, 2013. 2
- [60] Jiabei Zeng, Shiguang Shan, and Xilin Chen. Facial expression recognition with inconsistently annotated datasets. In *Proceedings of the European conference on computer vision (ECCV)*, pages 222–237, 2018. 6
- [61] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *AIES*, 2018. 3
- [62] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *CVPR*, 2021. 2
- [63] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *CVPR*, 2017. 2, 3