

Leveraging Diffusion Perturbations for Measuring Fairness in Computer Vision

Nicholas Lui^{1*}, Bryan Chia^{1*}, William Berrios², Candace Ross³, Douwe Kiela^{1,2}

¹Stanford University

²Contextual AI

³Meta AI

{niclui,bryancws}@stanford.edu

Abstract

Computer vision models have been known to encode harmful biases, leading to the potentially unfair treatment of historically marginalized groups, such as people of color. However, there remains a lack of datasets balanced along demographic traits that can be used to evaluate the downstream fairness of these models. In this work, we demonstrate that diffusion models can be leveraged to create such a dataset. We first use a diffusion model to generate a large set of images depicting various occupations. Subsequently, each image is edited using inpainting to generate multiple variants, where each variant refers to a different perceived race. Using this dataset, we benchmark several vision-language models on a multi-class occupation classification task. We find that images generated with non-Caucasian labels have a significantly higher occupation misclassification rate than images generated with Caucasian labels, and that several misclassifications are suggestive of racial biases. We measure a model’s downstream fairness by computing the standard deviation in the probability of predicting the true occupation label across the different perceived identity groups. Using this fairness metric, we find significant disparities between the evaluated vision-and-language models. We hope that our work demonstrates the potential value of diffusion methods for fairness evaluations.

Introduction

Computer vision systems have been shown to replicate harmful statistical associations found in the training data (Buolamwini and Gebru 2018; Stock and Cisse 2018; Wang et al. 2019). Encoding such biases increases the risk of computer vision models unfairly treating under-represented groups (Barocas et al. 2017), such as people of color (Buolamwini and Gebru 2018). To combat this, there have been several efforts to mitigate bias, ranging from sampling (Cao et al. 2020) to adversarial training (Alvi, Zisserman, and Nellåker 2018). However, there remains a lack of datasets, balanced along demographic traits, that are generally useful for evaluating the effectiveness of these techniques and the downstream fairness of computer vision models.

In this work, we explore the efficacy of diffusion methods (Sohl-Dickstein et al. 2015; Rombach et al. 2022) for generating datasets balanced along demographic traits that

can be used to evaluate the fairness of computer vision models. Our work is inspired by demographic perturbations in language (Maudslay et al. 2019; Smith and Williams 2021; Emmery et al. 2022; Qian et al. 2022), and extends it to images using recent advancements in diffusion models.

We propose measuring fairness via **diffusion perturbations**, a novel diffusion-based approach that can be used to generate a dataset balanced along demographic traits, such as the perceived race of people. In our fully automated approach, we use diffusion models to generate a large set of base images. Then, each image is edited using inpainting to generate multiple variants, where each variant refers to a different demographic group. Image sets that do not meet the bar for realism and prompt fidelity are filtered using a combination of a VQA and face attribution model.

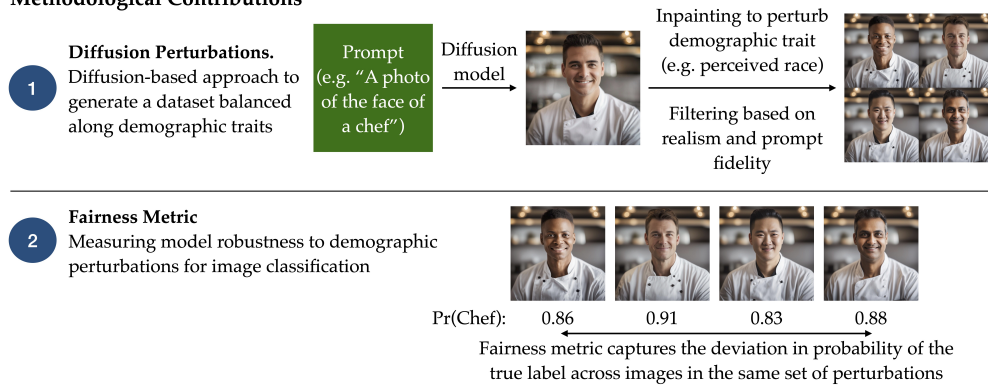
Images, in a given set of perturbations, share the same backgrounds and contexts, with the only difference being the perturbed demographic trait. To exploit the consistency within an image set, we propose a **fairness metric** that measures a model’s robustness to demographic perturbations. Given a classification task, the fairness metric measures the extent to which the probability of the true label varies across different demographic groups. We draw inspiration from perturbation metrics in language (Ma et al. 2021; Qian et al. 2022; Thrush et al. 2022) and extend them to images.

Using diffusion perturbations, we construct a **novel dataset depicting different occupations balanced by perceived race**. Each occupation comprises 1,200 images per perceived identity group with 4 identity groups represented (Black, Caucasian, East Asian, and Indian). We validate our dataset using crowdworkers, and show that our images demonstrate a high level of visual realism and a high probability of belonging to the intended target group.

Finally, we undertake a **fairness analysis using our occupations dataset**, where we evaluate several vision-language models (e.g. CLIP, FLAVA) using an occupation classification task. Using our fairness metric, we identify significant disparities between the models’ robustness to demographic perturbations. We also find that images generated with non-Caucasian labels have a lower classification accuracy than Caucasian-labeled images, and that many of these misclassifications are suggestive of model biases.

The paper’s main contributions are shown in Figure 1. To enable greater exploration of our work, we release our gen-

*These authors contributed equally.

Methodological Contributions**Dataset & Analysis Contributions**

- 3 Occupations Dataset** A novel, race-balanced dataset depicting various occupations
- 4 Occupations Fairness Analysis** Evaluating the downstream fairness of vision-language models on an occupation classification task using our dataset

Figure 1: Our main contributions. We propose (1) a novel diffusion-based approach to generate a dataset balanced along demographic traits, and (2) a fairness metric to measure a model’s robustness to demographic perturbations. We apply these techniques to (3) the creation of an occupations dataset and (4) produce fairness insights.

erated dataset at this link: bit.ly/occupation-dataset. We plan to release our code under a permissive license at this link: github.com/niclui/diffusion-perturbations.

Related Work

Demographic perturbations of images. To demographically perturb images, previous research has explored the use of generative adversarial networks (GANs) (Yucer et al. 2020; Dash, Balasubramanian, and Sharma 2022; Jain, Memon, and Togelius 2023), as well as reinforcement learning-based approaches (Wang and Deng 2020).

However, we favor a diffusion-based approach over a GAN-based one for two reasons. First, diffusion models have been shown to produce images with significantly more realism than GANs (Dhariwal and Nichol 2021) and are thus a more suitable choice for producing realistic perturbations. Second, to achieve the balance between realism and faithfulness to the user input, GANs often require additional training data or loss functions for individual applications (Meng et al. 2021). In contrast, our diffusion-based approach does not require task-specific training to produce images that display a high level of realism and faithfulness.

Datasets balanced along demographic traits. Of the few datasets that are balanced along demographic traits, most of them comprise close-up face images with self-reported ethnicity and demographic information. One example is the FairFace dataset (Karkkainen and Joo 2021).

We believe that our diffusion perturbations approach extends these datasets in two key ways. First, by generating a set of perturbations which share the same backgrounds and contexts, we are able to isolate the correlation between the demographic trait and model predictions. We do not have this consistency across images in real-world datasets.

Second, while existing datasets are valuable for evaluating bias in facial recognition, their utility in other downstream tasks may be limited. With control over the input prompts and thus the content of images generated, text-conditioned diffusion models can generate datasets that are tailored for fairness evaluations on a wider array of downstream tasks. We demonstrate this capability by generating a race-balanced dataset for occupation classification.

Nonetheless, we acknowledge that existing datasets, with self-reported demographic information, are more likely to contain diverse representations of people from different demographic groups. In contrast, the training data that diffusion models are trained on may contain more limited representations of different demographic groups.

Method

In this section, we describe how we use text prompts to generate images, filter images to ensure high quality, and generate masks to perturb the perceived demographic trait.

Prompt Creation

The first step is compiling a list of prompts that we use to generate images. We use the following 5 occupations as they are distributed across white and blue collar jobs. For occupations that are more difficult to generate, we include the distinguishing attire of the occupation to improve identifiability. Our prompts begin with “A photo of the face of”:

1. A car mechanic
2. A chef in a chef’s jacket
3. A commercial pilot
4. A doctor in a white coat with a stethoscope
5. A firefighter

We opt to generate a set of base images rather than use real images of people for two reasons. First, inpainting for perturbations generally performs better on a base synthetic image. Second, using a mix of real and synthetic images could potentially confound downstream evaluations.

Text-to-Image Generation

We use the Stable Diffusion model (Rombach et al. 2022), specifically Stable Diffusion XL for its improved photo-realism (Podell et al. 2023), to generate a 1024x1024 image for each prompt. In our initial experiments, we used an earlier version of Stable Diffusion (v2.1) with a grid search over a range of parameters to select the best image using a realism scorer. We find that images generated by Stable Diffusion XL, without any tuning, are significantly more realistic. We discuss our choice of hyperparameters in Appendix A1. For each occupation category, we generate 5-10k of images. We refer to these images as our base images.

Automated Filtering of Base Image Using VQA Model

We use a ViLT-B/32 VQA model (Kim, Son, and Kim 2021), fine-tuned on VQAv2 (Goyal et al. 2017), to evaluate the base images. We evaluate the following:

Text-to-image faithfulness. Following Hu et al. (2023), we use the VQA model to evaluate the faithfulness of the generated image to its text input. Specifically, we ask the VQA model “Is there a <occupation> in this image?” [Q1] where <occupation> is the occupation in the original prompt used to generate the image.

Limb realism. Diffusion models often face difficulties in producing realistic limbs. We thus ask the VQA model “Are this person’s limbs distorted?” [Q2].

Overall realism. The VQA model is asked: “Is this image real or fake?” [Q3].

Our filtering proceeds in two stages. First, we only keep base images which answer “Yes” to Q1 and “No” to Q2. We record the yield for each occupation in Appendix A2. For each occupation, we select the top ~2000 images with the highest score from the VQA model to Q3. From this set of ~2000 images, we filter images that are grayscale by computing the number of unique colors in an image. This gives us a set of base images which we will perturb.

Mask Generation

We pass each base image through an end-to-end segmentation pipeline. The pipeline comprises two state-of-the-art models: First, the base image is passed through an open-set object detection model, Grounded-DINO (Liu et al. 2023). Using the text prompt “person”, we obtain a bounding box around the person(s) in each image. Next, the base image and bounding box are passed into the Segment Anything model (Kirillov et al. 2023) which identifies a segmentation mask while conditioning on the input box.

Perturbation Using Inpainting

We seek to perturb the perceived race of the people in our images. To do so, we use the Stable Diffusion inpainting pipeline (Rombach et al. 2022). We pass a base image-mask pair into this pipeline, which inpaints the masked portion of the base image using the new prompt. The new prompt that we feed in is a perturbed version of the original prompt where we include a race identifier before the occupation (e.g. “A photo of the face of a [Black|Caucasian|Asian|Indian] firefighter”). The inpainting pipeline comprises 2 stages. First, the Stable Diffusion XL model (Podell et al. 2023) is used to generate latents of the desired output size. Second, a specialized high-resolution refinement model (Podell et al. 2023) applies the SDEdit image editing technique (Meng et al. 2021) to the latents generated in the first step, producing a high-quality edited image.

We use four race categories - Caucasian, Black, East Asian, Indian. Three widely studied identity groups in the AI fairness literature are Caucasian, Black, and Asian. However, given text prompts including “Asian”, images generated by Stable Diffusion are typically ones we perceive as East Asian.¹ To introduce greater diversity in our dataset, we additionally include one of the largest Asian groups that is not East Asian (Indian). We note that some real-world face image datasets, such as UTKFace (Zhang, Song, and Qi 2017), have used these exact same race labels.

Filtering Perturbed Sets Using FairFace

The perturbation process is not guaranteed to produce realistic representations of different identity groups. To mitigate this risk, we use the FairFace model. It is a race attribution model trained on the FairFace dataset, which is a novel face image dataset that is balanced on race (Karkkainen and Joo 2021). The FairFace model exhibits significantly higher accuracy when applied to novel face image datasets compared to models trained on imbalanced race datasets. Moreover, it maintains consistent accuracy across various race and gender groups. We use the 4-race version of the model which uses the exact same labels that we do (“Black”, “Caucasian”, “East Asian”, “Indian”). Using the FairFace model, we keep only sets of perturbed images where all 4 images are classified to the intended category. To maintain an even distribution, we sample 1200 image sets from each occupation. Our dataset comprises 24k images (1200×4 perceived races×5 occupations), with example images in Figure 2.

Limitations of Image Generation

Our approach yields edited images that are high quality and consistent with the base image. However, we note the following challenges: First, while we seek to generate images that are race-balanced, there is no guarantee that our images are balanced along other demographic traits. For instance, individuals in our base and edited images could be perceived as being more masculine. This could be attributed to bias in Stable Diffusion as many of our chosen occupations have a male skew (Luccioni et al. 2023). This issue limits the

¹As such, for East Asian-labeled images, we use “Asian” as the race identifier in the text prompt.



Figure 2: Samples of images generated for each occupation. The base image is generated using the text prompt “A photo of the face of a <occupation>”. We then generate a mask over the person(s) in the image. The original prompt is perturbed 4 times to include 4 different race identifiers: “A photo of the face of a [Black|Caucasian|Asian|Indian] <occupation>”. Base image-mask pairs are passed into the inpainting pipeline which produces four variants of the base image.

generalizability of our analysis to examining racial disparities within certain demographic boundaries (e.g. individuals who are perceived to be more masculine). To mitigate concerns about gender bias, we perform a robustness check using a regenerated sample of the dataset where we specify the perceived gender in each prompt.

Second, for occupations with a white skew in the US (e.g. pilot), we find that the base image tends to be Caucasian-presenting, possibly due to bias in Stable Diffusion. One question is whether the perceived race of the base image affects the quality of perturbations for other identity groups. However, since we use diffusion inpainting (Rombach et al. 2022; Wang et al. 2023), the individual in the base image is masked and it is less likely that their perceived race will have downstream effects on the perturbation quality. This belief is validated through our crowdworker review, which shows that there is no statistically significant difference in the image quality between the perceived identity groups.

Dataset Review

To validate the efficacy of the aforementioned pipeline, we conducted an Amazon MTurk survey over an evenly distributed sample of 4000 images (16.7% of the dataset). We asked workers two questions: (1) “Does this image contain obvious quality issues with the person? (e.g. blurred out facial features, additional limbs)”, with “Yes”, “No”,

and “Unsure” as possible responses, and (2) “What is one identity group that this person is likely to belong to?” with “Black”, “Caucasian”, “East Asian (e.g. Chinese)”, “South Asian (e.g. Indian)”, and “Others” as possible responses. We present our results in Table 1 and discuss further in the Appendix A3, where we show that the results are similar across the different perceived races.

Data	Realism Score	Race Fidelity Score
Overall	85.1%	91.0%
Chef	84.9%	90.6%
Doctor	86.6%	86.9%
Firefighter	91.5%	92.1%
Mechanic	77.4%	95.5%
Pilot	85.0%	90.0%

Table 1: Human dataset review: “Realism Score” is the proportion of reviewers who did not find quality issues with the image; “Race Fidelity Score” is the proportion of reviewers who indicated that the perceived race of the person is the same as the race specified during image editing.

Fairness Task

Our downstream task is multi-class occupation classification where the model chooses from a set of occupation labels.

Our analysis does not try to show that one model is generally more or less fair than another on occupation classification. Instead, we are evaluating the relative fairness of models for a specific set of occupations with a specific set of labels. We consider two label sets:

Base label set. We use the same set of labels for all occupations. The label set contains each original label as well as a negative label that is similar but clearly distinct from the original occupation. Labels used: [“chef”, “server”, “doctor”, “nurse”, “pilot”, “driver”, “mechanic”, “engineer”, “firefighter”, “police officer”].

Difficult label set. To provide a more challenging benchmark, we consider a more difficult set of occupation-specific labels as shown in Table 2. For each occupation, the label set includes the true occupation and 7 other adjacent occupations that we selected based on occupations listed in the U.S. Bureau of Labor Statistics. We include negative labels that are in the same line of work as the true label but with differing responsibilities (e.g. “doctor” and “nurse”).

A few of our negative labels may not be easily distinguishable in appearance from the true label (e.g. “doctor” and “physician assistant”). That being said, if the model is truly fair, we would expect it to choose the true label over the contending one with the same probability for all perceived identity groups.

To create a label set more appropriate for models trained on a larger amount of data, we use occupations that are less well-known, and thus better suited for models that have encountered them in their larger training datasets.

Occupation	Adjacent Occupations
Chef	Line cook, Cafeteria attendant, Waiter, Dishwasher, Food preparation worker, Host, Server
Doctor	Nurse, Physician assistant, Veterinarian, Clinical laboratory technician, Pharmacist, Emergency medical technician, Midwife
Firefighter	Fire chief, Coast guard, Security guard, Paramedic, Pilot, Police officer, Soldier
Mechanic	Automobile engineer, Civil engineer, Aerospace engineer, Mechanical engineer, Electrical engineer, Industrial engineer, Petroleum engineer
Pilot	Flight steward, Flight stewardess, Driver Aircraft fueler, Airline reservation agent, Air traffic controller, Aircraft engineer

Table 2: Difficult label set for multi-class classification.

Models Evaluated

Table 3 lists the evaluated models. We evaluate FLAVA (Singh et al. 2022) and unless otherwise stated, the ViT-B/32 variant of CLIP (Radford et al. 2021). For all models, we compute the cosine similarity between the image and each possible label which is prepended by “A photo

of”. The set of cosine similarities is passed into a softmax function to generate a set of prediction probabilities.

Model	Training Dataset	Dataset Size
FLAVA	PMD	70M
CLIP-OpenAI	WebImageText	400M
CLIP-LAION400M	Common Crawl	400M
CLIP-LAION2B	Common Crawl	2B

Table 3: Models evaluated on occupation classification task.

Fairness Metric

We define an extrinsic fairness metric to measure robustness to demographic perturbation in our classification task. If a model is fair, perturbing the demographic group should have minimal effect on model performance. We thus construct a fairness metric that captures the standard deviation in probability of the true label within an image set.

Formally, we have a classifier f_C and a dataset X . The dataset X comprises N image sets and each image set contains K perturbed images. For image j from set i , our classifier outputs the probability of the true label, $f_C(x_{ij})$. We compute the standard deviation of this probability across all j images to give us the standard deviation for set i . We then take the median standard deviation across all N sets.

Our fairness metric is 1 minus the median standard deviation. A fairness metric close to 1 (i.e. median standard deviation that is close to 0) implies that the model does equally well on every perceived identity group.

$$F_M(f_C, X) = 1 - \text{med}\left\{\sqrt{\frac{\sum_{j=1}^K (f_C(x_{ij}) - \overline{f_C(x_{ij})})^2}{K-1}}\right\}_{i=1, \dots, N}$$

Results

Results for Base Labels

Table 4 reports the results of the fairness metric across the different models. The models all achieve a very high level of parity across the different perceived identity groups. This necessitates a much harder set of labels to benchmark the models’ downstream fairness.

The table also shows the classification accuracy rate. However, we are primarily interested in the fairness metric. Regardless of what the average accuracy is, the model should be doing equally good/bad across the different perceived identity groups.

Model	Fairness Metric	Classification Accuracy
FLAVA	0.990	92.0%
CLIP-LAION400M	0.997	98.3%
CLIP-LAION2B	0.998	98.9%
CLIP-OpenAI	0.983	95.2%

Table 4: Models evaluated on our downstream occupation classification task using base label set.

Results for Difficult Labels

Table 5 reports the results of the fairness metric across the different models. We perform pairwise comparisons for the fairness metric using Mood’s non-parametric median test and Bonferroni’s Correction to account for multiple hypothesis testing. We find that FLAVA has the highest fairness metric at the 1% significance level. Within the CLIP models, we find that CLIP-OpenAI > CLIP-LAION2B > CLIP-LAION400M at the 1% level. The table also shows the accuracy rates for the different models.

Model	Fairness Metric	Classification Accuracy
FLAVA	0.983	90.1%
CLIP-LAION400M	0.835	75.1%
CLIP-LAION2B	0.849	79.8%
CLIP-OpenAI	0.884	68.6%

Table 5: Models evaluated on our downstream occupation classification task using difficult label set.

Although FLAVA has the smallest dataset size (trained on >5x less data than the CLIP models), it has the highest accuracy. Surprisingly, FLAVA achieves a similar accuracy on the difficult labels as it does on the base labels, while the other CLIP models all see a large decrease. FLAVA switches from being the least accurate model on the base labels to being the most accurate model on the difficult labels.

We hypothesize that because FLAVA is trained on a much smaller dataset than CLIP, it has a more limited vocabulary. Consequently, it might not be attempting to discern between a “doctor” vs “physician’s assistant” for instance, and is thus not tricked by the difficult label set. Thus, FLAVA outperforms all the CLIP models in both accuracy and the fairness metric. While CLIP’s richer model embedding space allows us to use more granular labels, it might also make it easier to draw out biases from the model.

However, this risk can be mitigated if the larger dataset used to develop richer embeddings is well curated. Within the CLIP models, CLIP-OpenAI scores the highest fairness metric. It achieves a significantly higher metric than CLIP-LAION400M, despite the two models having a similar architecture and training dataset size. The different dataset used could be responsible for disparities in the fairness metric. Birhane, Prabhu, and Kahembwe (2021) found that there were major data filtering issues with LAION-400M, resulting in the dataset containing large amounts of racist content and stereotypes. Thus, CLIP-LAION400M may have encoded racial biases to a larger extent than CLIP-OpenAI which was likely trained on a better curated dataset.

CLIP-LAION2B also has a lower fairness metric than CLIP-OpenAI, despite training on roughly 5x the amount of data. This could suggest that there are diminishing returns to fairness improvements with dataset size, especially if the underlying data contains significant biases.

Dataset robustness. In our dataset curation, we used only base images that a VQA model assessed to have included the target occupation. However, this might introduce biases in our evaluations in favor of models trained

on a similar dataset as the VQA model. Thus, we repeat the base experiment on a subset of the data where the base image was correctly classified by all the evaluated models. We have the same fairness metric ordering: FLAVA (0.98), CLIP-OpenAI (0.884), CLIP-LAION2B (0.861), CLIP-LAION400M (0.835). Full results are provided in the Appendix A4.

Label robustness. For a given occupation, we use 7 negative labels. However, some labels may be simply adding noise to the results. We remove every negative label except the top misclassified label and see if the results still hold. We retain the same ordering in the fairness metric with this experiment: FLAVA (0.999), CLIP-OpenAI (0.876), CLIP-LAION2B (0.866), CLIP-LAION400M (0.847). Full results are provided in the Appendix A5.

Ablation study on number of parameters. We use the ViT-B/32 variant of CLIP, but it’s not clear if our results hold for a larger parameter size. We evaluate the ViT-L/14 variant of CLIP, which has close to 3X the number of parameters. We find that CLIP-LAION400M (ViT-L/14) still has the lowest fairness metric, while CLIP-OpenAI (ViT-L/14) and CLIP-LAION2B (ViT-L/14) have similar fairness metrics. Our results are in Table 6. We continue to use the ViT-B/32 variant for our analysis.

Model	Fairness Metric	Classification Accuracy
CLIP-LAION400M	0.862	82.4%
CLIP-LAION2B	0.908	90.1%
CLIP-OpenAI	0.903	84.1%

Table 6: Ablation Study. We evaluate the ViT-L/14 variant of CLIP on our downstream occupation classification task.

Robustness to gender bias in image generation and perturbation. As part of the limitations of image generation, we noted that the occupations we evaluate may have a male skew, resulting in generated images that are largely male-presenting. In our error analysis (Appendix A9), we also find that the presenting gender in the inpainted image may sometimes differ from the original base image.

To ensure that our results are robust to these biases, we regenerate a small sample of our dataset, while ensuring that it is balanced across both the male and female genders. We do so by specifying a gender in the prompt used for image generation and inpainting (e.g. “A photo of the face of a [male|female] firefighter”). Given that we specify the same gender in the prompts used for base image generation and inpainting, it is much less likely that the presenting gender in the inpainted image will differ from that of the original base image. For each occupation, our sample comprises 200 images (100 male, 100 female) for each perceived race.

Using this sample dataset, we find that our fairness metric ordering is retained (results in parentheses): FLAVA (0.944) > CLIP-OpenAI (0.874) > CLIP-LAION2B (0.827) > CLIP-LAION400M (0.821). These results suggest that our qualitative results are robust to gender bias in image generation and perturbation.

Occupation-Level Analyses

Across the board, all 3 non-Caucasian perceived identity groups have a lower classification accuracy than Caucasians at the 1% level (Black: -6.09%, East Asian: -3.21%, Indian: -3.20%). Table 7 shows the results of the fairness metric for each occupation. Within CLIP models, CLIP-LAION2B achieves the highest fairness metric on Chef, while CLIP-OpenAI does the best on every other occupation.

Model	Chef	Doc	FF	Mec	Pilot
FLAVA	0.999	0.728	0.988	0.986	0.973
CLIP LAION400M	0.837	0.830	0.946	0.791	0.747
CLIP LAION2B	0.956	0.740	0.931	0.798	0.749
CLIP OpenAI	0.895	0.867	0.955	0.816	0.859

Table 7: Fairness Metric by Occupation. Abbreviations: Doc: Doctor, FF: Firefighter, Mec: Mechanic

In the Appendix A7, we illustrate how different models perform when predicting the occupation of the image across different perceived identity groups. To get race-level effects, we regress a binary outcome variable (1 if image predicted correctly, 0 otherwise) on perceived race, using robust standard errors clustered at the image set level to account for within-set correlation.

We show findings from 3 occupations below, but analyze all 5 occupations in the Appendix A7.

Chef: CLIP-LAION400M and CLIP-OpenAI predict a Black chef with an accuracy 10-13% lower than Caucasians. CLIP-LAION400M predicts an East Asian and Indian Chef with a 10% and 9% lower accuracy respectively. This disparity can be explained by CLIP-LAION400M and CLIP-OpenAI predicting non-Caucasian chefs as line cooks more often than Caucasians. Summing across both models, Black chefs are predicted as line cooks 104% more times than Caucasians. East Asians and Indians are predicted as line cooks 72% and 57% more times respectively.

Doctor: The top misclassified label is physician assistant. This is understandable as a physician assistant and doctor has similar appearances. Still, it does not nullify the fact that we should see an equal probability of predicting either class for the different perceived identity groups. Instead, we find a significantly lower difference in predicting a Black doctor than Caucasians (ranging from 10 – 24%) across the CLIP-based models. Black doctors are misclassified as physician assistants 34% more times than Caucasians. For CLIP-LAION400M, Indians and East Asians have a 12% *higher* accuracy than Caucasians. In a cursory probe, we find that there is a corresponding pattern in classifying Caucasian-presenting doctors within CLIP-LAION2B’s training set as doctors more often than Black-presenting doctors, who are classified as physician assistants. We provide further details in Appendix A8.

Pilot: We find that the probability of CLIP models predicting Black, Indian, and East Asian pilots are significantly lower as compared to Caucasian pilots, with most of these models predicting “flight stewards” instead. Given that 95.7% of employed pilots are Caucasian (according to the US Bureau of Labor Statistics), it is possible that there is a lack of diversity in representations of pilot in the US-centric portions of the training data.

Error analysis. For every non-Caucasian perceived identity group, we sample 10 image sets where the Caucasian-labeled image is correctly classified and the non-Caucasian group is not. We present the images for each occupation in the Appendix A9. For most sets, we find that the misclassified images look similar to the correctly classified Caucasian-labeled image, suggesting that it may be model bias, rather than image quality, that is driving inequitable model performance.

Implicit Association Test (IAT)

Drawing from the racial IAT (Greenwald, McGhee, and Schwartz 1998; Maina et al. 2018) to analyze subvert biases, we probe the model to choose one of two labels: “A trustworthy/untrustworthy person”. We report results in Appendix A6. We find that CLIP-LAION400M has the largest spread in probabilities across the perceived race groups, consistent with our finding that CLIP-LAION400M could be less fair than the other two CLIP models in some situations.

Evaluation Using Large Language Models

Finally, we are interested in assessing the fairness of Large Language Models using difficult labels in our dataset. For this, we employed LENS (Berrios et al. 2023), a model that reasons and performs classification tasks based on visual descriptions. We use the prompt “Question: Can you please identify the occupation that best represents the image? Short Answer: {answer}” to select the best continuation. Since LENS provides the log probabilities for each possible continuation, we calculate the joint probability and then employ the same calculation as in CLIP or FLAVA.

In our case, LENS (H_{14} -FlanT5_{xxl}) achieved a fairness metric of 0.837, while LENS (H_{14} -FlanT5_{xl}) reached 0.965. These results suggest that LENS (H_{14} -FlanT5_{xl}) may be performing more fairly on our task than the CLIP models. Finally, Table 8 shows occupation-specific fairness metric results across both LENS models.

Model (LENS)	Chef	Doc	FF	Mec	Pilot
H_{14} -FlanT5 _{xxl}	0.928	0.826	0.824	0.762	0.916
H_{14} -FlanT5 _{xl}	0.984	0.894	0.978	0.939	0.986

Table 8: Fairness Metric by Occupation for LENS models. Doc: Doctor, FF: Firefighter, Mec: Mechanic

Conclusion

As computer vision models grow larger and more powerful, there is an increasing need to develop robust techniques to audit their fairness on downstream tasks. We propose

a novel diffusion-based approach to generate datasets balanced along demographic traits that can be used to assess relative model fairness. In our approach, a large set of images is generated using diffusion models and then perturbed to generate multiple variants representing different demographic groups. We develop a dataset on occupation classification, and show that our approach can be used to expose model biases. We also develop a fairness metric to measure a model’s robustness to demographic perturbations. Using our fairness metric and generated dataset, we uncover disparities between several vision-language models. We hope that our work demonstrates the value of diffusion models to fairness evaluations and research.

Resources

The Appendix can be found at this link: bit.ly/dp-appendix. To enable greater exploration and improvement on our work, we release our dataset here: bit.ly/occupation-dataset. We plan to release our code under a permissive license at this link: github.com/niclui/diffusion-perturbations.

Ethical Statement

While our approach shows promise in building a dataset balanced along demographic traits for fairness evaluations, we recognize that there are key ethical considerations that need to be addressed.

First, we acknowledge that race is socially constructed and that it is impossible to identify someone’s race solely based on their appearance. Racial identity is a personal decision and there is no fixed definition of what people of a specific race should look like. Nonetheless, we recognize that a lot of the harms from racism come from racial profiling (Glover 2009), which is heavily influenced by a person’s appearance and the extent to which the individual has physical characteristics that societal perceptions associate with their race (Maddox and Perry 2018). Despite the limitations of Stable Diffusion, we believe that the generated images contain some of those characteristics. Nonetheless, we acknowledge that our work may have the unintended consequence of reifying a fixed definition of race.

Second, diffusion models may not generate images that reflect the diversity of people belonging to different identity groups. Image generations build on the representations of different identity groups found in the training dataset, which may be especially narrowly defined for people of color. In future work, we hope to evaluate methods like FairDiffusion (Friedrich et al. 2023) which can instruct generative models to produce images that are fairer and more diverse. We discuss other limitations of our image generation process in the section detailing our methods.

Acknowledgements

We would like to thank Helen Gu, Amir Hertz, Gautam Mittal, Rajan Vivek, Rohan Koodli, and Melissa Hall for their thoughtful feedback. This work was funded in part by a cloud compute grant from Google and the Stanford Institute for Human Centered Artificial Intelligence. NL is supported by a Stanford Knight-Hennessy Fellowship. Meta was not

involved in running any models for the dataset creation or evaluation.

References

- Alvi, M.; Zisserman, A.; and Nellåker, C. 2018. Turning a blind eye: Explicit removal of biases and variation from deep neural network embeddings. In *Proceedings of the European Conference on Computer Vision Workshops*.
- Barocas, S.; Crawford, K.; Shapiro, A.; and Wallach, H. 2017. The Problem with Bias: From Allocative to Representational Harms in Machine Learning. In *Special Interest Group for Computing, Information and Society (SIGCIS)*.
- Berrios, W.; Mittal, G.; Thrush, T.; Kiela, D.; and Singh, A. 2023. Towards Language Models That Can See: Computer Vision Through the LENS of Natural Language. *arXiv:2306.16410*.
- Birhane, A.; Prabhu, V. U.; and Kahembwe, E. 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963*.
- Buolamwini, J.; and Gebru, T. 2018. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of Machine Learning Research*.
- Cao, D.; Zhu, X.; Huang, X.; Guo, J.; and Lei, Z. 2020. Domain balancing: Face recognition on long-tailed domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Dash, S.; Balasubramanian, V. N.; and Sharma, A. 2022. Evaluating and mitigating bias in image classifiers: A causal perspective using counterfactuals. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. In *Proceedings of Advances in Neural Information Processing Systems*.
- Emmery, C.; Kádár, Á.; Chrupała, G.; and Daelemans, W. 2022. Cyberbullying classifiers are sensitive to model-agnostic perturbations. *arXiv preprint arXiv:2201.06384*.
- Friedrich, F.; Schramowski, P.; Brack, M.; Struppek, L.; Hintersdorf, D.; Luccioni, S.; and Kersting, K. 2023. Fair diffusion: Instructing text-to-image generation models on fairness. *arXiv preprint arXiv:2302.10893*.
- Glover, K. S. 2009. *Racial profiling: Research, racism, and resistance*. Rowman & Littlefield Publishers.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Greenwald, A. G.; McGhee, D. E.; and Schwartz, J. L. 1998. Measuring individual differences in implicit cognition: the implicit association test. *Journal of personality and social psychology*.
- Hu, Y.; Liu, B.; Kasai, J.; Wang, Y.; Ostendorf, M.; Krishna, R.; and Smith, N. A. 2023. TIFA: Accurate and interpretable text-to-image faithfulness evaluation with question answering. *arXiv preprint arXiv:2303.11897*.

- Jain, A.; Memon, N.; and Togelius, J. 2023. Zero-shot racially balanced dataset generation using an existing biased StyleGAN2. *arXiv preprint arXiv:2305.07710*.
- Karkkainen, K.; and Joo, J. 2021. Fairface: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*.
- Kim, W.; Son, B.; and Kim, I. 2021. ViLT: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment anything. *arXiv preprint arXiv:2304.02643*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; and Zhang, L. 2023. Grounding DINO: Marrying DINO with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.
- Luccioni, A. S.; Akiki, C.; Mitchell, M.; and Jernite, Y. 2023. Stable bias: Analyzing societal representations in diffusion models. *arXiv preprint arXiv:2303.11408*.
- Ma, Z.; Ethayarajh, K.; Thrush, T.; Jain, S.; Wu, L.; Jia, R.; Potts, C.; Williams, A.; and Kiela, D. 2021. Dynaboard: An evaluation-as-a-service platform for holistic next-generation benchmarking. In *Proceedings of Advances in Neural Information Processing Systems*.
- Maddox, K. B.; and Perry, J. M. 2018. Racial appearance bias: Improving evidence-based policies to address racial disparities. *Policy Insights from the Behavioral and Brain Sciences*.
- Maina, I. W.; Belton, T. D.; Ginzberg, S.; Singh, A.; and Johnson, T. J. 2018. A decade of studying implicit racial/ethnic bias in healthcare providers using the implicit association test. *Social science & medicine*.
- Maudslay, R. H.; Gonen, H.; Cotterell, R.; and Teufel, S. 2019. It's all in the name: Mitigating gender bias with name-based counterfactual data substitution. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2023. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. *arXiv preprint arXiv:2307.01952*.
- Qian, R.; Ross, C.; Fernandes, J.; Smith, E.; Kiela, D.; and Williams, A. 2022. Perturbation augmentation for fairer nlp. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Singh, A.; Hu, R.; Goswami, V.; Couairon, G.; Galuba, W.; Rohrbach, M.; and Kiela, D. 2022. FLAVA: A foundational language and vision alignment model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Smith, E. M.; and Williams, A. 2021. Hi, my name is Martha: Using names to measure and mitigate bias in generative dialogue models. *arXiv preprint arXiv:2109.03300*.
- Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*.
- Stock, P.; and Cisse, M. 2018. Convnets and imagenet beyond accuracy: Understanding mistakes and uncovering biases. In *Proceedings of the European Conference on Computer Vision*.
- Thrush, T.; Tirumala, K.; Gupta, A.; Bartolo, M.; Rodriguez, P.; Kane, T.; Rojas, W. G.; Mattson, P.; Williams, A.; and Kiela, D. 2022. Dynatask: A framework for creating dynamic AI benchmark tasks. In *Proceedings of the Association for Computational Linguistics: System Demonstrations*.
- Wang, M.; and Deng, W. 2020. Mitigating bias in face recognition using skewness-aware reinforcement learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wang, S.; Saharia, C.; Montgomery, C.; Pont-Tuset, J.; Noy, S.; Pellegrini, S.; Onoe, Y.; Laszlo, S.; Fleet, D. J.; Soricut, R.; et al. 2023. Imagen editor and editbench: Advancing and evaluating text-guided image inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wang, T.; Zhao, J.; Yatskar, M.; Chang, K.-W.; and Ordonez, V. 2019. Balanced Datasets Are Not Enough: Estimating and Mitigating Gender Bias in Deep Image Representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Yücer, S.; Akçay, S.; Al-Moubayed, N.; and Breckon, T. P. 2020. Exploring racial bias within face recognition via per-subject adversarially-enabled data augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*.
- Zhang, Z.; Song, Y.; and Qi, H. 2017. Age Progression/Regression by Conditional Adversarial Autoencoder. In *IEEE Conference on Computer Vision and Pattern Recognition*.