

Fully Unsupervised Self-debiasing of Text-to-Image Diffusion Models

Korada Sri Vardhana Shrikrishna Lolla Soma Biswas

Indian Institute of Science, Bengaluru

srivardhanak@iisc.ac.in shrikrishnalolla@gmail.com somabiswas@iisc.ac.in

Abstract

*Text-to-image (T2I) diffusion models have achieved widespread success due to their ability to generate high-resolution, photorealistic images. These models are trained on large-scale datasets, like LAION-5B, often scraped from the internet. However, since this data contains numerous biases, the models inherently learn and reproduce them, resulting in stereotypical outputs. We introduce **SelfDebias**, a fully unsupervised test-time debiasing method applicable to any diffusion model that uses a UNet as its noise predictor. SelfDebias identifies semantic clusters in an image encoder’s embedding space and uses these clusters to guide the diffusion process during inference, minimizing the KL divergence between the output distribution and the uniform distribution. Unlike supervised approaches, SelfDebias does not require human-annotated datasets or external classifiers trained for each generated concept. Instead, it is designed to automatically identify semantic modes. Extensive experiments show that SelfDebias generalizes across prompts and diffusion model architectures, including both conditional and unconditional models. It not only effectively debiases images along key demographic dimensions while maintaining the visual fidelity of the generated images, but also more abstract concepts for which identifying biases is also challenging.*

1. Introduction

Generative models [10, 12, 17], particularly T2I models [27], have achieved huge success in the past few years due to their ability to generate high-resolution, photorealistic images that align well with human perceptions. These models are often trained on massive datasets [31], which are scraped from the internet. Since a significant amount of bias exists in internet-sourced content, these models unintentionally learn such biases. Recent studies [36] have shown that these biases are not only learned by the models but are also reproduced and sometimes amplified in their outputs. Since diffusion models are now commonly used in real-world systems [25, 28] and for data augmentation in

downstream tasks [38], the bias in the model’s outputs, if left unchecked, can lead to unfair outcomes and significant representational harms [23].

A natural solution to prevent this bias in the model’s output is to re-train the model on curated unbiased datasets. But this approach suffers from degradation of diversity and visual fidelity [25], in addition to the cost required to curate such datasets and retrain the models. As a result, several alternative debiasing strategies have been proposed [6, 9, 15, 23]. Majority of the existing supervised approaches primarily address societal bias like gender and race, which are well known and can be pre-defined. Thus, these specific biased concepts can be learnt, and further used for debiasing. But, supervised debiasing frameworks are quite restrictive for free-form image generation: (i) Learning each concept separately, usually by generating sufficient number of examples of each using the model can be prohibitively expensive when debiasing along multiple combinations of attributes [9, 23]; (ii) In several scenarios, manually specifying the bias is very difficult, if not impossible. For example, for abstract concepts like a peaceful scenery, concrete concepts like bags, etc., it is not clear what are the biases introduced! Thus, it is imperative to develop unsupervised techniques which can automatically discover the inherent biases and self-correct itself to generate more unbiased outputs. This is an extremely challenging and relatively unexplored direction. Recently, unsupervised frameworks such as TIME [22] and UCE [9] have been proposed. However, both methods rely on a predefined list of attributes and their combinations that are intended to be debiased. Unlike TIME, which is restricted to editing attributes with only two categories, UCE enables the simultaneous debiasing of multiple categorical attributes. Consequently, most of our comparisons in this paper are made with UCE. UCE modifies the influence a token exerts along attribute directions in a specified ratio, making prior knowledge of these attributes essential, which may not always hold true in practice.

Towards this goal, we propose *SelfDebias*, a fully unsupervised debiasing framework that operates directly in the semantic latent space (h -space) of diffusion models. Rather than relying on predefined bias axes, we use the CLIP [26]

image encoder to uncover semantic latent modes of variation through unsupervised clustering of generated image embeddings. These clusters serve as *proxy* semantic groups, and we guide the generation process to reduce the dominance of any single group. Although we use CLIP in our experiments, any image encoder capable of semantically separating images in its embedding space could be used. This utilization of semantic separability in the embedding space, which is already available, enables us to perform debiasing.

The proposed SelfDebias framework has three main modules: **(i) semantic projection module:** Here, a large batch of images is generated using a conditional diffusion model based on a given prompt. Then a projection network is learnt, which maps the h -space activations from the UNet of these generated images to their semantic image embeddings; The **(ii) semantic mode discovery** module performs hierarchical clustering to automatically cluster the projected features and discover semantic modes in the representation space; During inference time, the **(iii) self-debias module** uses these learned modes to perform soft cluster assignment over the projected h -space activations. The KL divergence between the predicted soft cluster distribution and a uniform prior acts as a differentiable surrogate for balancing the concepts. The resulting loss is back-propagated to edit the h -space representations mid-generation, and the updated activations are used in the subsequent denoising steps. This process allows us to apply semantic guidance directly in h -space without modifying the model weights or relying on any external supervision, making our method lightweight and compatible with existing pretrained diffusion models.

Terminology. In the LLM literature [30], “self-debiasing” often denotes reflective or prompt-based procedures that leverage a model’s own knowledge to critique or revise its outputs. In this paper, by *SelfDebias* we mean a purely *test-time*, *statistical* intervention on hidden activations (the h -space). Crucially, model weights (U-Net, VAE, text encoder) are *never* updated.

Our key contributions are as follows:

1. To the best of our knowledge, SelfDebias is the first **fully** unsupervised, test-time framework for debiasing diffusion models, without requiring attribute labels, external classifiers, or retraining.
2. We demonstrate that our inference-time intervention-based framework generalizes across different demographic attributes like gender, race, age, etc. without any supervision.
3. We showcase the debiasing capability of SelfDebias in scenarios where the appropriate debiasing axis is unclear or difficult to define, such as abstract prompts, which cannot be handled by supervised techniques.

2. Related Works

Diffusion-based models are now the state-of-the-art for high-fidelity image generation. However, a growing body of research has shown that these models are not free from societal biases, like gender-occupation stereotypes, race, and age, [1, 4, 21, 21, 32]. The existing debiasing techniques can broadly be categorized into two approaches, namely debiasing with or without requiring retraining.

Debiasing with retraining. Many prior works attempt to debias generative models by retraining from scratch, as seen in the progression of Stable Diffusion (e.g., SD 1.5 vs. SD 2.1), or through fine-tuning [5, 37, 39]. Some approaches use a balanced reference dataset to perform bias correction. Choi et al. [5] proposed a density-ratio estimation technique for bias identification and reweighted the training data accordingly. More recently, in the context of diffusion models, Shen et al. [34] propose a distributional alignment loss that directly optimizes for fairness in the generated outputs. Few recent works [13, 14, 16] have also explored debiasing by either fine-tuning or retraining the model on carefully curated datasets. While these retraining-based approaches have shown promising results, they depend on labeled or curated datasets and involves a substantial computational overhead for the retraining procedure. In contrast, our method does not require any retraining, access to the original training data, or demographic labels, substantially lowering the cost and complexity of deployment in real-world settings.

Debiasing without retraining. These methods aim to intervene during inference time to mitigate bias, as opposed to fine-tuning-based approaches. As a result, they are significantly more cost-efficient in terms of training. A recent work by Parihar et al. [23] proposes a method for debiasing diffusion models by editing the intermediate latent representations (h -space) of the U-Net during inference. Similarly [18, 20] directly perform edits in h -space to control semantic attributes, but both of these require at least some form of demographic supervision or annotation. Chuang et al. [6] performs algebraic manipulations of biased directions in text embeddings to debias $T2I$ diffusion models. Jiang et al. [15] argue that since text-image alignment occurs primarily through cross-attention, updating these layers using parameter-efficient tuning (e.g., LoRA-style adaptation) enables controlled attribute steering. [9, 22] follows same principles and targets the cross-attention layers to debias and erase at the same time. A recent work [35] detects bias features embedded in diffusion models at neuron level and intervenes to prevent bias. However, all these approaches assume access to labeled data and demographic dimensions to supervise optimization. NoiseCLR [7] also discovers self-supervised directions, but it requires costly training in noise space, does not apply to unconditional generation, and its directions capture general semantic varia-

tions rather than bias-specific axes. *This limits their general applicability, especially in cases where it is difficult to define debiasing axes. In contrast, we propose a fully unsupervised self-debiasing framework, which is applicable for a much wider range of concepts.*

3. Proposed Framework

Here, we describe SelfDebias, a fully unsupervised debiasing framework for pretrained diffusion models, without retraining, supervision, or access to attribute labels. The key idea is to align the distribution of generated samples across latent semantic modes, discovered via clustering in the image encoder’s embedding space, to a uniform target, thereby reducing the inherent bias. The framework consists of three key modules (Fig. 1): (1) Semantic Projection Module: where a projection from h -space to semantic space is learnt; (2) Semantic Mode Discovery Module: where semantic clusters are determined via unsupervised clustering; and (3) Self-debias Module: where h -space editing via a KL-divergence loss is performed between the predicted cluster distribution and a uniform prior. We will now describe these modules.

3.1. Semantic Projection Module

Here, we learn a projection from the h -space to a semantically disentangled latent space.

Intermediate Activations (h -space): We briefly describe the h -space of diffusion models, which is the intermediate representation for completion. The U-Net backbone used in diffusion models has a bottleneck feature layer between the encoder and decoder blocks. The activations at this bottleneck layer are referred to as h -space [18], denoted h_t . These intermediate representations contain rich semantic information about the current sample at timestep t , and have been shown to encode interpretable features such as gender, age, and object type [18]. In our framework, we extract h_t during the sampling process and apply gradient-based updates. This makes our method efficient, modular, and applicable to any diffusion model with U-Net-style architectures.

Recent supervised debiasing framework [23] has utilized h -space. First, we perform a simple analysis on the h -space to understand whether they are sufficiently discriminative and can be directly clustered to determine the underlying concepts. We take gender (male and female) as the running example, for ease of understanding. We generated 2,000 images using the prompt "A face of a person" and capture their h -space embeddings during generation. These embeddings are clustered and visualized, and the corresponding images are also labeled using CLIP as a zero-shot classifier. Both visualizations are shown side by side in Fig. 2. We observe that although these features are semantically rich, they are also highly entangled. This prevents effective clustering in this space for the purpose of extracting distinct semantic

modes. Therefore, we need to project the h -space vectors into an embedding space where the representations are more semantically disentangled.

Let $h_t^{(i)} \in \mathbb{R}^d$ denote the h -space activation at timestep t for image i , extracted from the bottleneck layer of the U-Net. Let $x_0^{(i)}$ be the corresponding generated image obtained via decoding, and let $z^{(i)} = E(x_0^{(i)})$ be its embedding in the semantic space, where E is the image encoder (we used CLIP), capable of encoding images into a semantically disentangled space. We train a lightweight projection network g_ψ to align h -space activations with their corresponding semantic embeddings. To achieve this, we use the NT-Xent contrastive loss [2], which pulls together positive pairs $(g_\psi(h_t^{(i)}), z^{(i)})$ while pushing apart negative pairs $(g_\psi(h_t^{(i)}), z^{(j)})$ for $j \neq i$ within the same batch. For a batch of size N , the loss for a sample i is defined as:

$$\mathcal{L}_{\text{proj}}^{(i)} = -\log \frac{\exp(\cos(g_\psi(h_t^{(i)}), z^{(i)})/\tau)}{\sum_{j=1}^N \exp(\cos(g_\psi(h_t^{(i)}), z^{(j)})/\tau)}, \quad (1)$$

where $\cos(\cdot, \cdot)$ denotes cosine similarity and τ is the temperature hyperparameter. The total projection loss over the batch is given by:

$$\mathcal{L}_{\text{proj}} = \frac{1}{N} \sum_{i=1}^N \mathcal{L}_{\text{proj}}^{(i)}. \quad (2)$$

This encourages the projected h -space features to align with their corresponding semantic embeddings, while remaining well-separated from others in the batch. Once trained, g_ψ is frozen and used during inference to project h -space activations into the semantic embedding space.

3.2. Semantic Mode Discovery Module

After semantic projection, we obtain a set of projected semantic embeddings $\{\hat{z}_t^{(i)}\}_{i=1}^N$ of N images, where $\hat{z}_t^{(i)} = g_\psi(h_t^{(i)})$. To uncover latent semantic modes of variation without supervision, we perform clustering on the projected embeddings. We do not assume prior knowledge of the number of clusters, as the relevant modes of variation may not be obvious.

Challenge in discovering modes: To better understand what kind of clustering is required to compute the different concepts, we again consider the widely used demographic concepts, but here, instead of only gender, we also consider other factors like race, age, etc. We observe that though gender is well clustered, the other factors are quite entangled. As expected, we cannot disentangle across multiple dimensions like race and gender simultaneously, since each gender will have people of all races. With more dimensions, the entanglement further increases. To discover all these

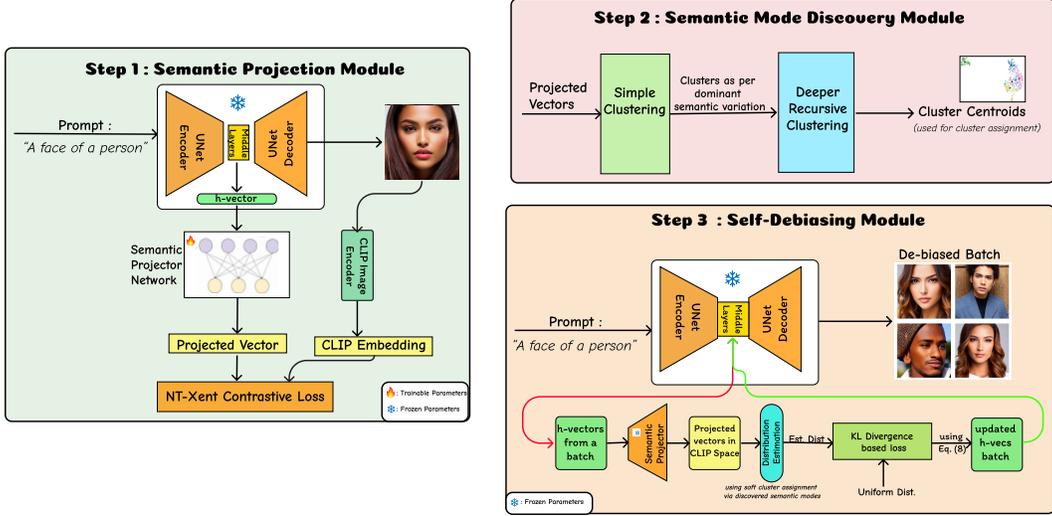


Figure 1. An overview of the proposed SelfDebias framework. 1) First, the projection network is trained to transform the semantically entangled h -space vectors to a disentangled semantic space; 2) Second, the semantic modes are discovered through a two stage clustering process of the projected h -vectors; 3) Finally, the h -vectors of a biased batch are updated such that the tend towards a more uniform distribution, leading to a debiased batch.

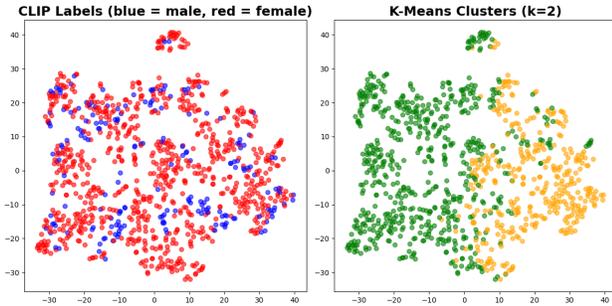


Figure 2. Left: h -space vectors labeled using CLIP as a zero-shot classifier (blue = male, red = female). Right: The same h -vectors grouped via K-Means clustering ($k = 2$). Both visualizations are produced using PCA followed by t-SNE. The unsupervised clustering (right) fails to align meaningfully with the semantic labels from CLIP (left). This misalignment highlights the entangled nature of the h -space.

factors in an unsupervised manner is extremely challenging. Here, we propose to use a hierarchical clustering [29] with adaptive weighting to discover these semantic models to eventually debias the model output.

First Stage Clustering: To estimate the appropriate number of clusters, we evaluate clustering quality across different k values using silhouette score [33], which quantifies how well each point fits within its assigned cluster relative to other clusters. For each k , we compute the average silhouette score across all data points. Higher score indicates that the clusters are well-separated and internally cohesive. We select the k value closest to the maximum silhouette

score. This aligns with the usual model selection principles in unsupervised learning.

Second Stage Clustering: The clusters obtained up to this stage capture the dominant semantic groupings in the data, i.e., those aligned along axes of highest semantic variation. However, each cluster may still contain finer-grained semantic distinctions. We apply recursive spectral clustering within each to uncover finer semantic dimensions. This recursive process is governed by two hyperparameters: s_{\min} (minimum allowed cluster size) and d_{\max} (maximum recursion depth). A cluster is split only if it contains more than s_{\min} samples and its current depth $d < d_{\max}$. This results in a hierarchy of clusters with varying semantic granularity. We denote the final set of leaf clusters as $\mathcal{C} = \{c_1, c_2, \dots, c_K\}$, where each cluster c_j is associated with a semantic embedding space centroid and a depth $d_j \in \{0, 1, \dots, d_{\max}\}$. For each projected vector $\hat{z}_t^{(i)}$, we compute a soft assignment over the final set of clusters using cosine similarity with temperature scaling:

$$p_t^{(i)}[j] = \frac{\exp(\alpha \cdot \cos(\hat{z}_t^{(i)}, c_j))}{\sum_{k'=1}^K \exp(\alpha \cdot \cos(\hat{z}_t^{(i)}, c_{k'}))} \quad (3)$$

where α is a temperature parameter controlling assignment sharpness. We then average these soft assignments across the images to obtain the empirical distribution over discovered clusters as follows:

$$P_t[j] = \frac{1}{N} \sum_{i=1}^N p_t^{(i)}[j] \quad (4)$$

Adaptive Cluster Weightage: To reflect the hierarchical structure of the clusters, we assign each cluster c_j a weight based on its depth:

$$w_j = \frac{1}{d_j + 1} \quad (5)$$

The target distribution is then defined as:

$$\mathcal{U}[j] = \frac{w_j}{\sum_{k=1}^K w_k} \quad (6)$$

This depth-weighted uniform target ensures that shallower semantic modes are treated as more significant, while allowing fine-grained sub-clusters to contribute proportionally.

3.3. Self-debiasing Module

To encourage balanced representation of the discovered latent modes, inspired by [23], we minimize the KL divergence between the predicted distribution P_t and the depth-weighted uniform target \mathcal{U} :

$$\mathcal{L}_{\text{uniform}} = \text{KL}(P_t \parallel \mathcal{U}) = \sum_{j=1}^K P_t[j] \log \left(\frac{P_t[j]}{\mathcal{U}[j]} \right) \quad (7)$$

Our novel formulation facilitates debiasing across all the axes (gender, age, race) simultaneously. Thus the target distribution will have all the concepts that have been discovered in the previous module. Since the assignment probabilities are defined via a softmax over cosine similarities, they form a differentiable function of the projected h -vectors, making the loss in Eq. (7) differentiable. We back-propagate this loss into the h -space representations:

$$\tilde{h}_t^{(i)} = h_t^{(i)} - \gamma \cdot \nabla_{h_t^{(i)}} \mathcal{L}_{\text{uniform}} \quad (8)$$

The updated $\tilde{h}_t^{(i)}$ is then used to compute the denoised latent $x_{t-1}^{(i)}$ in the sampling trajectory. Notably, the entire process is unsupervised and test-time only, requiring no retraining, no attribute labels, and no access to the training dataset.

Overall Framework: To summarize, we begin by generating a set of 2,000 images along with their corresponding h -space activations for a given prompt. These activations are then used to train the semantic projection module. Once trained, this module is used to project the h -space vectors into the semantic embedding space, where clustering is performed. After identifying the cluster centroids, we intervene in the h -space during inference, using the method described above, to steer each generated image batch towards more balanced outputs.

4. Experiments

Here, we present extensive experiments to evaluate the efficacy of the proposed framework. First, we evaluate the

effectiveness of our approach in debiasing the face domain, addressing both major semantic shifts and finer-grained variations, and compare it against existing methods. Additionally, we assess SelfDebias on occupational bias using prompts derived from the WinoBias [42] dataset. Notably, we highlight SelfDebias’s unique ability to debias even when the target attributes are vague or unspecified. We further demonstrate the method’s independence from any specific image encoder and its model-agnostic nature.

4.1. Implementation Details

For our experiments, we use Stable Diffusion v1.5 [27]. All experiments are conducted on NVIDIA RTX A5000 GPUs. We used a batch size of 100 for dataset generation and performed image synthesis with 50 inference steps. For training the semantic projection module, we use a time-conditioned two-layer MLP projecting to 512 dimensions. Its input is a flattened h -vector and a timestep embedding and it outputs a 512-dimensional vector aligned with the CLIP space. For evaluation, we use Fairness Discrepancy (FD) [5] to quantify demographic bias, which measures the deviation of the generated attribute distribution from a uniform target. FD is computed as:

$$\text{FD} = \left\| \bar{p} - \mathbb{E}_{x \sim p_\theta(x)} [C(x)] \right\|_2, \quad (9)$$

where $p_\theta(x)$ is the distribution of generated images and \bar{p} denotes the uniform distribution over attribute classes. $C(\cdot)$ is a high-accuracy attribute classifier, for which we have used same setup as in [23]. We used FID [11] to determine the quality of the images generated. We have used standard implementation of clean-fid [24] for calculating it.

Baseline Approaches: To the best of our knowledge, there are no fully unsupervised techniques, so we compare with the following approaches: 1) **Supervised Approaches:** Here, we compare with i) Random Sampling, which is vanilla stable diffusion model, without any modifications. ii) ITI-Gen [41], which works through appending of prompt tokens with learned prompt embeddings of each category. iii) Fair Diffusion [8], which uses a lookup table for recognizing bias from the given text input and adds corresponding scaled embeddings to the text embeddings iv) H-Guidance [23], which not only assumes the knowledge of the bias concepts, but also learns the specific concepts by generating sufficient images using the base diffusion model; 2) **Unsupervised Approach with Bias Knowledge:** We also compare with the recent unsupervised approach UCE [9] and self-supervised approach ID [20], both of which assume the knowledge of the bias concepts for performing debiasing, but doesn’t use any labeled datasets to achieve their debiasing objectives. Unlike our method, ID [20] learns a separate direction for each attribute; for fairness, we selected concepts via a uniform-probability switch rather than manual toggling.

Table 1. Comparison of different methods for debiasing gender using FD and FID metrics. We also indicate whether the method assumes known bias labels or operates in an unsupervised manner. Underline = Best overall, **bold** = Best among unsupervised methods; Note that ours is completely unsupervised, whereas UCE and ID assume the knowledge of the bias axes.

Method	FD↓	FID↓	Bias unknown?	No supervision?
Random Sampling	0.317	72.37	—	—
ITI-Gen [41]	0.049	<u>64.79</u>	X	X
Fair Diffusion [8]	0.227	71.22	X	X
H-Guidance [23]	0.024	70.69	X	X
UCE [9]	0.028	69.81	X	✓
ID [20]	0.019	70.64	X	✓
SelfDebias (ours)	0.009	70.52	✓	✓

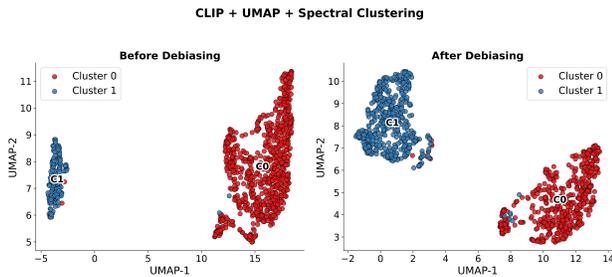


Figure 3. UMAP visualization of image embeddings before (left) and after (right) debiasing for gender. Right plot exhibits more balanced cluster distribution, indicating effective mitigation of bias.

4.2. Mitigating Bias in Human Face Generation

Here, we present the results of our approach using (i) two clusters and (ii) multiple clusters.

Two Cluster Results: We observe that the image distributions produced by diffusion models naturally separate into a few dominant semantic clusters in CLIP embedding space, which typically align with high-variance attributes such as gender in faces or species in pets. In the first experiment, we cluster the embeddings of human faces into two clusters (which naturally align with gender) and report the performance in Table 1. The results of the competing methods are taken directly from Parihar et al. [23] for consistency and comparability. We observe that despite being a completely unsupervised approach, it outperforms supervised techniques in terms of FD and is third in terms of FID (only slightly worse than UCE).

For qualitative analysis, we present visualizations of images in CLIP embedding space before and after debiasing in Fig. 3 and image grids before and after debiasing in Fig. 4. Together, these visualizations illustrate how our method restructures the semantic space and promotes greater balance across dominant demographics.

Multiple Clusters: Here, we further divide each main cluster found in above clustering for debiasing along more fine-

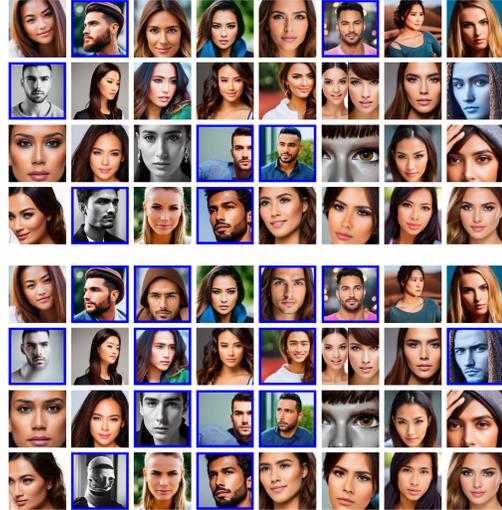


Figure 4. Top: Generated images using vanilla Stable Diffusion 1.5. Bottom: Generated images using SelfDebias. The debiased set contains 44% men, as opposed to 22% using the original model (in blue).

grained attributes, like age and race. Despite the fact that these sub-clusters cannot simultaneously align cleanly with all attributes, they often represent consistent local semantics, what we call *pseudo-semantic regions*. We observe that it captures local structure that reflects entangled but meaningful groupings.

We identify such regions recursively and use their centroids to guide debiasing. At inference time, we compute similarity of latents to these centroids and enforce a form of semantic diversity by encouraging a uniform sampling over them. Importantly, we do not attempt to identify each attribute axis explicitly, nor to enumerate all possible demographic intersections, both of which can be combinatorially large and often unobservable. Instead, we rely on the emergent structure revealed through recursive unsupervised clustering. Though the number of resulting clusters is often fewer than the total number of demographic combinations, we empirically find that enforcing balance across these pseudo-semantic regions reduces demographic bias, as measured by the evaluation metrics.

As shown in Table 2, recursive debiasing improves parity across gender, race and their combinations. Fig. 5 visualizes the discovered cluster hierarchy using UMAP, illustrating how deeper recursive levels uncover nuanced structure otherwise hidden in flat representations. We show sample images from each cluster in Fig. 6. We observe that SelfDebias outperforms even the supervised approaches for gender and Race+Age+Gender criteria, despite being a completely unsupervised approach, unlike the other compared approaches. *For these results, the supervised techniques generate images for each concept at a time, for example,*



Figure 5. UMAP projection showing fine-grained clusters and their centroids (marked with black-crosses) for faces generation.



Figure 6. Results of hierarchical clustering; each column displays a subset of images from one cluster. Column 1 primarily contains men. Cluster 3 includes individuals with darker skin tones, while Cluster 4 predominantly features individuals with Asian facial features. Clusters 2 and 5 appear more semantically mixed, including faces with characteristics commonly associated with Indian and white individuals.

for age, images are generated specifically for diverse ages, whereas for UCE [9], we enumerate all the concepts. Our framework debiases across all the dimensions at once.

4.3. Occupation-Related Bias

Next, we evaluate SelfDebias on a set of occupation-based prompts inspired by the Winobias benchmark (e.g., “a photo of a doctor”, “a photo of a teacher”), which are known to elicit biased gender associations in generative models due to correlations in web-scale training data. For each prompt, we generate a batch of images using our method and then classify the gender of each generated image to compute the deviation ratio [9] for evaluation. Importantly, we perform debiasing using cluster centroids computed from a single generic prompt: “a photo of a person”. This decouples the clustering step from the downstream prompts and demonstrates that SelfDebias does not require task-specific repre-

Table 2. Comparison of different methods for debiasing using the FD and FID metrics. CNA = Checkpoint Not Available. Criteria RAG = Race + Age + Gender. Race includes White, Black, Indian & Asian. Age includes Young, Old & Adult. Underline = Best overall, **bold** = Best among unsupervised methods, * = Uses supervision, \$ = Unsupervised but knows labels

Criteria	Random Sampling		H-Guid* [23]		UCE\$ [9]		SelfDebias (Ours)	
	FD	FID	FD	FID	FD	FID	FD	FID
Race	0.387	72.37	<u>0.188</u>	75.52	0.254	83.84	0.237	87.08
Age	0.488	72.37	<u>0.194</u>	78.88	0.238	87.82	0.272	87.08
Gender	0.317	72.37	0.024	70.69	0.028	69.81	0.015	87.08
RAG	0.385	72.37	CNA	CNA	0.314	89.96	0.297	87.08

sentations at test time. As shown in Table 3, our method consistently reduces gender skew, even though the semantic clusters were learned using a neutral base prompt. We observe that for all three of the four prompts, it outperforms UCE, whereas for Analyst, it performs almost at par with the others. The results for other methods in Table 3 are directly taken from [9] for consistency in comparison.

4.4. Debiasing Abstract Concepts

The proposed SelfDebias is unique in its ability to perform debiasing without any prior knowledge of the bias concepts. This enables our method to generalize beyond predefined notions of bias and operate effectively in settings where the source of bias is unknown or unlabelled. To demonstrate our method’s ability to operate without prior bias axes, we present two illustrative examples. In the first case, we prompt a Stable Diffusion model with the phrase “a peaceful moment”. The model typically generates either empty scenic landscapes or landscapes with people, often favoring the former disproportionately. In this context, it is not feasible to define an explicit bias axis, as the notion of “peacefulness” is abstract and subjective. Despite the lack of a clearly identifiable bias direction, SelfDebias effectively identifies and mitigates the over-representation of people in the outputs. As shown in Fig. 7, SelfDebias successfully debiases such ambiguous scenarios without relying on explicit guidance or external classifiers.

In the second example, we examine the model’s behavior when prompted with “a fantasy creature”. Interestingly, the outputs tend to cluster into two broad semantic categories: one consisting of humanoid entities and the other dominated by beast-like forms resembling dinosaurs. We observe a bias toward the latter, with the model disproportionately producing dragon-like creatures while under-representing humanoid variants. SelfDebias successfully identifies and mitigates this imbalance. As illustrated in Fig. 8, SelfDebias leads to a more semantically balanced distribution without requiring predefined concept axes.

Table 3. Profession-specific gender deviation ratios (mean \pm std) across different debiasing methods. Lower is better. Underline = Best overall, **bold** = Best among unsupervised methods

Profession	Original-SD	Concept Algebra [40]	Debias-VL [6]	TIME [22]	TIME + Preserve [9]	UCE [9]	SelfDebias
Librarian	0.86 \pm 0.06	0.66 \pm 0.07	0.34 \pm 0.06	0.26 \pm 0.05	0.35 \pm 0.01	0.07 \pm 0.07	0.06 \pm 0.05
Teacher	0.42 \pm 0.01	0.46 \pm 0.00	0.11 \pm 0.05	0.34 \pm 0.06	0.07 \pm 0.06	0.06 \pm 0.02	0.04 \pm 0.02
Analyst	0.58 \pm 0.12	0.24 \pm 0.18	0.71 \pm 0.02	0.52 \pm 0.03	0.13 \pm 0.05	0.20 \pm 0.07	0.22 \pm 0.06
Doctor	0.78 \pm 0.04	0.40 \pm 0.02	0.50 \pm 0.04	0.58 \pm 0.03	0.41 \pm 0.08	0.20 \pm 0.02	0.18 \pm 0.04

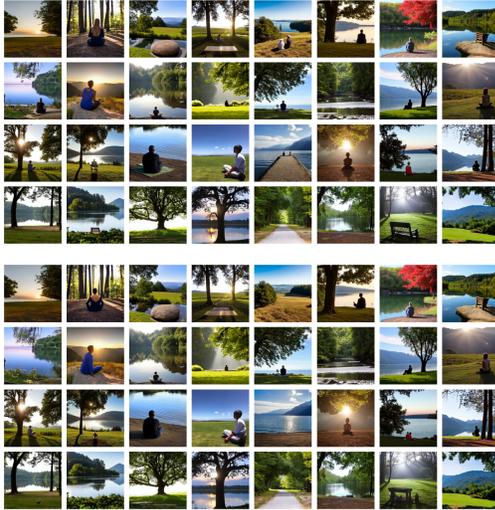


Figure 7. Top: Generated images using original model Bottom: Generated images after applying SelfDebias. In the original set (top), 20 out of 32 images (62.5%) depict scenes without people. After debiasing (bottom), this number decreases to 16 out of 32 images (50%), reflecting a 12.5 % reduction in people-less scenes.

4.5. Generalization

Generalization to Imbalanced Targets Our method generalizes to arbitrary target distributions over discovered semantic clusters. This is especially useful for simulating specific environments, such as demographic skews in urban settings. Fig. 2 in supplementary material shows results.

Generalization across image encoders: To evaluate the independence of our debiasing method from the choice of image encoder, we additionally experiment with OpenCLIP’s image encoder [3] for debiasing the gender attribute in the facial domain. As shown in Fig. 3 of supplementary material, our method successfully debiases the representations when using OpenCLIP.

Generalization across model To demonstrate the model-agnostic nature of SelfDebias, we evaluate it on an unconditional DDIM model trained on CelebA-HQ dataset [19]. Despite the absence of text conditioning, our framework reduces demographic imbalance. We have shown qualitative results in Fig. 4 of supplementary. This confirms that our method generalizes effectively across architectures and

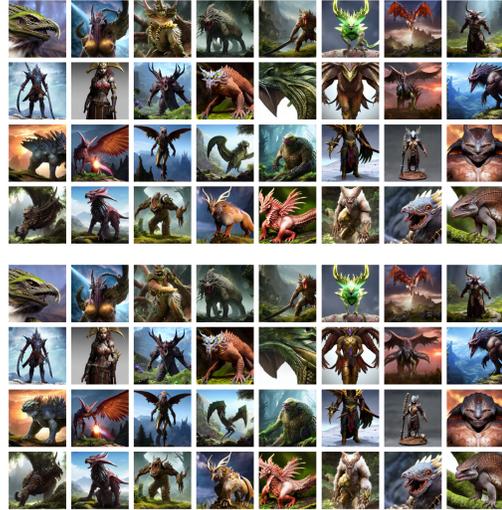


Figure 8. Top: Generated images using original model. Bottom: Generated images after applying SelfDebias. In the original set (top), only 30% of the scenes depict humanoid-like creatures, whereas in the debiased set (bottom), this increases to 40%, indicating a more balanced semantic distribution.

training regimes. Unsupervised methods like UCE [9] and TIME [22], cannot debias unconditional diffusion models because of their reliance on text tokens, whereas our method is generalizable to any diffusion model with an UNet.

5. Summary and Conclusion

We propose SelfDebias, a fully unsupervised, framework for mitigating bias in diffusion models. SelfDebias guides generation by promoting uniformity over semantic clusters discovered in a latent space, without relying on attribute labels or external classifiers. Operating entirely in the h -space, it generalizes across architectures that uses UNet as their backbone. A limitation of our method is that discovered clusters may not always align with human-interpretable concepts. It can successfully debias only if those concepts are generated by the model. In conclusion, we hope that SelfDebias contributes to the pursuit of fair generation, especially in settings where supervision is limited but ethical generation is critical, and serves as a strong baseline for future unsupervised debiasing methods.

Acknowledgment

This work is partially funded by the Kotak IISc AI-ML Centre, IISc, whose support is gratefully acknowledged.

References

- [1] Abhipsa Basu, R Venkatesh Babu, and Danish Pruthi. Inspecting the geographical representativeness of images from text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5136–5147, 2023. 2
- [2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmLR, 2020. 3
- [3] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2818–2829, 2023. 8
- [4] Jaemin Cho, Abhay Zala, and Mohit Bansal. Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3043–3054, 2023. 2
- [5] Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, and Stefano Ermon. Fair generative modeling via weak supervision. In *International Conference on Machine Learning*, pages 1887–1898. PMLR, 2020. 2, 5
- [6] Ching-Yao Chuang, Varun Jampani, Yuanzhen Li, Antonio Torralba, and Stefanie Jegelka. Debiasing vision-language models via biased prompts. *arXiv preprint arXiv:2302.00070*, 2023. 1, 2, 8
- [7] Yusuf Dalva and Pinar Yanardag. Noiseclr: A contrastive learning approach for unsupervised discovery of interpretable directions in diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24209–24218, 2024. 2
- [8] Felix Friedrich, Manuel Brack, Lukas Struppek, Dominik Hintersdorf, Patrick Schramowski, Sasha Luccioni, and Kristian Kersting. Fair diffusion: Instructing text-to-image generation models on fairness. *arXiv preprint arXiv:2302.10893*, 2023. 5, 6
- [9] Rohit Gandikota, Hadas Orgad, Yonatan Belinkov, Joanna Materzyńska, and David Bau. Unified concept editing in diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5111–5120, 2024. 1, 2, 5, 6, 7, 8
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 1
- [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 5
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1
- [13] Min Hou, Yueying Wu, Chang Xu, Yu-Hao Huang, Chenxi Bai, Le Wu, and Jiang Bian. Invdiff: Invariant guidance for bias mitigation in diffusion models, 2024. 2
- [14] Lin-Chun Huang, Ching Chieh Tsao, Fang-Yi Su, and Jung-Hsien Chiang. Debiasing diffusion model: Enhancing fairness through latent representation learning in stable diffusion model. *arXiv preprint arXiv:2503.12536*, 2025. 2
- [15] Yilei Jiang, Wei-Hong Li, Yiyuan Zhang, Minghong Cai, and Xiangyu Yue. Debiasing text-to-image diffusion models with self-discovering latent directions. 1, 2
- [16] Yeongmin Kim, Byeonghu Na, Minsang Park, JoonHo Jang, Dongjun Kim, Wanmo Kang, and Il-Chul Moon. Training unbiased diffusion models from biased dataset. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [17] Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013. 1
- [18] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. *arXiv preprint arXiv:2210.10960*, 2022. 2, 3
- [19] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5549–5558, 2020. 8
- [20] Hang Li, Chengzhi Shen, Philip Torr, Volker Tresp, and Jindong Gu. Self-discovering interpretable diffusion latent directions for responsible text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12006–12016, 2024. 2, 5, 6
- [21] Alexandra Sasha Luccioni, Christopher Akiki, Margaret Mitchell, and Yacine Jernite. Stable bias: Analyzing societal representations in diffusion models. *arXiv preprint arXiv:2303.11408*, 2023. 2
- [22] Hadas Orgad, Bahjat Kawar, and Yonatan Belinkov. Editing implicit assumptions in text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7053–7061, 2023. 1, 2, 8
- [23] Rishabh Parihar, Abhijnya Bhat, Abhipsa Basu, Saswat Mallick, Jogendra Nath Kundu, and R Venkatesh Babu. Balancing act: distribution-guided debiasing in diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6668–6678, 2024. 1, 2, 3, 5, 6, 7
- [24] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11410–11420, 2022. 5
- [25] Malsha V Perera and Vishal M Patel. Analyzing bias in diffusion-based face generation models. In *2023 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–10. IEEE, 2023. 1

- [26] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 1
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 5
- [28] Harrison Rosenberg, Shima Ahmed, Guruprasad Ramesh, Kassem Fawaz, and Ramya Korlakai Vinayak. Limitations of face image generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 14838–14846, 2024. 1
- [29] Maurice Roux. A comparative study of divisive and agglomerative hierarchical clustering algorithms. *Journal of Classification*, 35(2):345–366, 2018. 4
- [30] Timo Schick, Sahana Udupa, and Hinrich Schütze. Self-diagnosis and self-debiasing: A proposal for reducing corpus-based bias in nlp, 2021. 2
- [31] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models, 2022. 1
- [32] Preethi Seshadri, Sameer Singh, and Yanai Elazar. The bias amplification paradox in text-to-image generation. *arXiv preprint arXiv:2308.00755*, 2023. 2
- [33] Ketan Rajshekhar Shahapure and Charles Nicholas. Cluster quality analysis using silhouette score. In *2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 747–748, 2020. 4
- [34] Xudong Shen, Chao Du, Tianyu Pang, Min Lin, Yongkang Wong, and Mohan Kankanhalli. Finetuning text-to-image diffusion models for fairness. *arXiv preprint arXiv:2311.07604*, 2023. 2
- [35] Yingdong Shi, Changming Li, Yifan Wang, Yongxiang Zhao, Anqi Pang, Sibe Yang, Jingyi Yu, and Kan Ren. Dissecting and mitigating diffusion bias via mechanistic interpretability. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 8192–8202, 2025. 2
- [36] Yingdong Shi, Changming Li, Yifan Wang, Yongxiang Zhao, Anqi Pang, Sibe Yang, Jingyi Yu, and Kan Ren. Dissecting and mitigating diffusion bias via mechanistic interpretability. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8192–8202, 2025. 1
- [37] Christopher TH Teo, Milad Abdollahzadeh, and Ngai-Man Cheung. Fair generative models via transfer learning. In *Proceedings of the AAAI conference on artificial intelligence*, pages 2429–2437, 2023. 2
- [38] Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective data augmentation with diffusion models. *arXiv preprint arXiv:2302.07944*, 2023. 1
- [39] Soobin Um and Changho Suh. A fair generative model using lecam divergence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 10034–10042, 2023. 2
- [40] Zihao Wang, Lin Gui, Jeffrey Negrea, and Victor Veitch. Concept algebra for (score-based) text-controlled generative models. *Advances in Neural Information Processing Systems*, 36:35331–35349, 2023. 8
- [41] Cheng Zhang, Xuanbai Chen, Siqi Chai, Chen Henry Wu, Dmitry Lagun, Thabo Beeler, and Fernando De la Torre. Itigen: Inclusive text-to-image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3969–3980, 2023. 5, 6
- [42] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Gender bias in coreference resolution: Evaluation and debiasing methods. *arXiv preprint arXiv:1804.06876*, 2018. 5