

Generated Bias: Auditing Internal Bias Dynamics of Text-To-Image Generative Models

Abhishek Mandal¹, Susan Leavy², and Suzanne Little¹

¹ Dublin City University, Dublin, Ireland

² University College Dublin, Dublin, Ireland

abhishek.mandal2@mail.dcu.ie

susan.leavy@ucd.ie, suzanne.little@dcu.ie

Abstract. Text-To-Image (TTI) Diffusion Models such as DALL-E and Stable Diffusion are capable of generating images from text prompts. However, they have been shown to perpetuate gender stereotypes. These models process data internally in multiple stages and employ several constituent models, often trained separately. In this paper, we propose two novel metrics to measure bias internally in these multistage multimodal models. Diffusion Bias was developed to detect and measures bias introduced by the diffusion stage of the models. Bias Amplification measures amplification of bias during the text-to-image conversion process. Our experiments reveal that TTI models amplify gender bias, the diffusion process itself contributes to bias and that Stable Diffusion v2 is more prone to gender bias than DALL-E 2.

Keywords: Gender Bias · Generative Computer Vision · Diffusion Models

1 Introduction

Text-to-image (TTI) Diffusion Models such as Stable Diffusion and DALL-E have impressive capabilities of generating images from text descriptions. However, they can learn bias inherited from training data including gender bias. This presents as stereotypical results, highly imbalanced class distributions or low visual diversity of outputs. Much research has shown that models such as DALL-E and Stable Diffusion perpetuate stereotypical gender bias such as generating more images of men for *engineers* and images of women for *nurses* [4, 8, 10]. Quantitative measures such as the Multimodal Composite Association Score (MCAS) [10] have shown the extent of stereotypical gender representations in both DALL-E and Stable Diffusion [8]. This issue also affects vision-language foundation models such as Contrastive Language Image Pretraining (CLIP), which is used by both these diffusion models for embedding input text prompts [12]. Bias in generative vision models, coupled with their growing popularity and widespread use, have the potential to cause significant social harm, perpetuating gender inequality and undermining social progress. Therefore, it is important to understand, detect, and measure gender bias in these models.

Previous research auditing bias in TTI models has focussed on these models as a whole (as black-boxes) [3,4,8,13,18]. TTI models are multistage multimodal models often employing different models processing data internally in multiple stages and modalities [14,15]. Often the constituent models are trained separately. Metrics for measuring and analysing how bias is handled internally by deep learning models has largely been limited to simpler models like image classifiers [17,21]. Therefore, it is important to develop metrics that can detect and measure internal bias in TTI models which can provide greater insight into the internal bias dynamics and the role of the model architecture and help better understand and mitigate bias in these models. We recognise that gender is complex and multi-faceted and while we do not seek to reinforce a binary view of gender, this research considers men and women in a binary sense with a view to building on this to include multiple genders in future research.

We hypothesise that bias is amplified by the TTI models and that the model architecture plays an important role in this. To investigate, we propose two novel metrics: **Diffusion Bias** (δ), which measures bias contributed by the diffusion process and **Bias Amplification** (α), which measures bias amplification by the model during conversion from text to image. We seek to answer the following research questions:

1. How can we effectively measure internal bias in TTI models?
2. Is bias amplified during the image generation process and how does the model architecture affect it?

To summarise, the contributions from this paper are:

1. Introduction of two novel metrics, based on the Multimodal Composite Association Score, to detect and measure gender bias in TTI diffusion models.
2. Analysis of the internal bias dynamics of TTI models and the role of the model architecture in bias amplification.

2 Background and Related Work

Generative Text-To-Image models are complex, multistage and multimodal models that are often composed of multiple separate models. The constituent models are often pre-trained, as in the use of CLIP in DALL-E and Stable Diffusion, and the final network is then retrained [14,15]. The result of this complexity makes accurate and independent bias analytics challenging [8,10]. In this section we describe the main components of the key TTI diffusion models, possible causes of bias and summarise the current efforts to understand, measure and mitigate gender bias in such models.

2.1 Text-To-Image (TTI) Diffusion Models

TTI diffusion models generally employ a three-stage mechanism to generate images from input text called prompts. First, the input prompt describing the

desired output image is converted into embeddings using a multimodal vision-language model. The embeddings are processed and prepared for the diffusion part to create priors. Finally, the priors are fed to the diffusion model to generate the output image [14, 15].

DALL-E 2 [14] is the second version of DALL-E and was released in 2022. It has two components: CLIP and unCLIP. In the first part, it uses CLIP to generate embeddings of the text inputted by the user. The embeddings are then modelled using a Gaussian diffusion model and is called the *Diffusion Prior*. The diffusion prior is continuous and consists of a decoder-only Transformer with a causal attention mask on a sequence containing: the encoded text/caption; the CLIP text embeddings; diffusion timestep encoding; the noised CLIP embedding of the image; and a final embedding from the Transformer, in that order. The diffusion model (unCLIP) then models from the representation space (from the prior) to generate the image via reverse diffusion.

Stable Diffusion v2.0 [15] was released in 2021 by LMU Munich and Runway ML. It is based on a new type of diffusion model called a *Latent Diffusion Model* (LDM). It is a multi-stage multimodal model similar to DALL-E and uses CLIP for the initial text encoding. It consists of three main components: CLIP for text encoding, a UNet+ scheduler for gradual diffusion from the latent space and an autoencoder decoder for the final image generation. The diffusion process is run in the ‘latent space’ to conserve computing resources and boost speed.

2.2 Sources of Bias in TTI Diffusion Models

Gender bias in computer vision has been examined from various perspectives including the source of bias [20, 21], contributing factors and attempts to measure and/or reduce bias [17, 19].

One probable source of bias in TTI models, as in other machine learning models, is the training data, which is often scraped from the Internet. The data present on the Internet mirrors the biases of our society and these are reflected in the resulting models [1, 6, 10]. Bias embedded in training data can then be amplified internally by a deep learning model [8, 9]. Mandal *et al.* [9] showed that model architecture also plays an important role in how bias is handled by a vision model with Vision Transformers learning biased representations more than Convolutional Neural Networks.

Biases from data or other components are passed down the training pipeline and, critically, amplified by various factors such as the component model architecture and the training methodology [9, 17]. Recently Friedrich *et al.* [5] also demonstrated the effect of multilinguality on gender stereotype magnification. Previous methodologies for analysing internal bias amplification in vision models have been limited to simpler classification models such as CNNs [17, 21].

Large and complex models such as TTI diffusion models using multiple different models may internally exhibit bias in different ways. DALL-E 2 and Stable Diffusion v2.0 both use CLIP for initial prompt embeddings and are trained on similar data [14, 15] but show different levels of apparent gender bias when subjected to the same prompts [10]. To better understand this behaviour, we

propose methodologies and metrics for measuring bias amplification in TTI diffusion models.

2.3 Evaluating Gender Bias in TTI Models

The complexity of TTI Diffusion models makes the identification, interpretation, measurement and mitigation of potential bias very challenging. Identification can occur through benchmarking efforts, such as Jha *et al.* [7] who recently extended an existing textual dataset with visual depictions of identity groups to demonstrate a pull towards stereotypical depictions in nationality-based biases.

Bias is often studied intersectionally. Luccioni *et al.* [8] studied the presence of intersectional gender bias in Text-to-Image (TTI) models by evaluating their output using image captioning models and creating clusters based on visual features. Their tool, StableBias, also allowed for visual analysis of the outputs. They used prompts which included multiple identities such as occupation, ethnicity, and gender. Their tool allows for exploratory analysis of the output of TTI models but does not allow for quantitative measurement of bias, especially in the representation space.

Similarly, TIBET proposed by Chinchure *et al.* [3], measures bias along multiple axes such as physical appearance, ableism, gender, religion, and race. They used a VQA similar to [8] to extract concepts from images generated by TTI models. They used counterfactuals to analyse how much the bias changed along the axes. The use of an external VQA model can be an issue here as it may introduce or reflect the bias of the VQA model.

A more intersectional approach is taken by Cho *et al.* [4] where gender and skin tone is evaluated in images generated by popular TTI models using both human and automated methods. They found Stable Diffusion to generate more images of a specific skin tone or gender than DALL-E. The authors used exploratory analysis of the outputs.

Vice *et al.* [18] uses three metrics for quantifying bias in TTI models: *Distribution bias* which measures the distribution of bias in the TTI generated output, and their novel metrics; *Jaccard Hallucination* which measures the correlation of hallucinations and bias; and *Generative Miss Rate* which measures how bias affects model performance. These metrics measure bias in a post hoc way similar to the metrics discussed above.

The approaches so far are mostly post hoc analyses of the output but offer limited insight into how the bias is handled internally by the TTI models. In previous work (Mandal *et al.* [10]), we studied the presence of gender bias in DALL-E 2 and Stable Diffusion v2 and found significant gender bias over different categories. Both models are more likely to generate a greater number of male-looking images for traditionally male-dominated occupations such as *CEO*, *engineer*, *doctor*, and *programmer* and a greater number of female-looking images for traditionally female-dominated occupations such as *beautician*, *nurse*, *librarian*, and *housekeeper*. We proposed *MCAS* to measure stereotypical gender associations in the internal representation of TTI models. MCAS is a linear scale with the sign indicating the nature and the value indicating the magnitude of

the bias. *Diffusion Bias* and *Bias Amplification* are based on this concept and are defined in Section 3.

We identified two major gaps in the previous research. (1) Evaluation via exploratory analysis of the results for bias analytics, mostly using human or other captioning or VQA models. This can lead to human or the captioning model’s bias influencing the analysis. This also limits how bias is represented and handled in the internal representation space of the TTI models. (2) The methods and metrics measure bias for the whole model. Thus it is difficult to understand how bias is generated, amplified or mitigated internally. We address these gaps in our work.

3 Multimodal Composite Association Score (MCAS)

The Multimodal Composite Association Score (MCAS) [10] was developed from the Word Embeddings Association Test (WEAT) [2], which itself is based on the popular Implicit Association Test. MCAS consists of four components each of which measures the relative association of real-world target concepts such as occupations and sports to male and female genders in four different combinations of visual and textual modalities. It measures the stereotypical association of real-world concepts (e.g., occupations and sports) called *Targets* and gender concepts (e.g., male and female) called *Attributes*. Both the targets and attributes consist of images and texts representing each concept (Tables 1, 2). All the images are generated using prompts.

Let A and B be two sets of gender attributes and W be a set of targets (e.g. occupation). Then

$$s(w, A, B) = \text{mean}_{a \in A} \cos(\mathbf{w}, \mathbf{a}) - \text{mean}_{b \in B} \cos(\mathbf{w}, \mathbf{b}) \quad (1)$$

where, $w \in W$, $\cos(w, a)$ and $\cos(w, b)$ denote the cosine similarities between the embedding vectors from the sets W and A and B respectively, and $s(w, A, B)$ is the association score between w , A , and B . Suppose the target, w , is closer to A . In that case, this will result in a positive association score and meaning it is biased towards A and if it is closer to B , the score will be negative indicating bias towards B . MCAS is composed of four such scores, each calculated as:

$$\text{AssociationScore} = \text{mean}_{w \in W} s(w, A, B) \quad (2)$$

The constituent association scores are:

- **Image-Image Association Score (II_{AS}):** measures bias in visual modality between image attributes representing gender and generated images representing target concepts.
- **Image-Text Prompt Association Score (ITP_{AS}):** is a visual-linguistic score measuring bias between the image attributes representing gender and the textual prompts used to generate the target concepts.

- **Image-Text Attributes Association Score (IT_{AS})**: is also a visual-linguistic score which measures bias similar to the other scores with the difference being that the attributes are represented not by images, but by text.
- **Text-Text Association Score (TT_{AS})**: measures bias in the textual modality with the attributes being the same as IT_{AS} and the targets being the same as ITP_{AS} . This is the only score which is entirely textual and as both DALL-E 2 and Stable Diffusion use CLIP for text embedding, this score also measures CLIP bias.

MCAS is calculated as the sum of all the four association scores as:

$$MCAS = II_{AS} + ITP_{AS} + ITA_{AS} + TT_{AS} \quad (3)$$

MCAS measures bias in the internal embeddings of the models and not against any external benchmark. Figure 1 shows the components of MCAS.

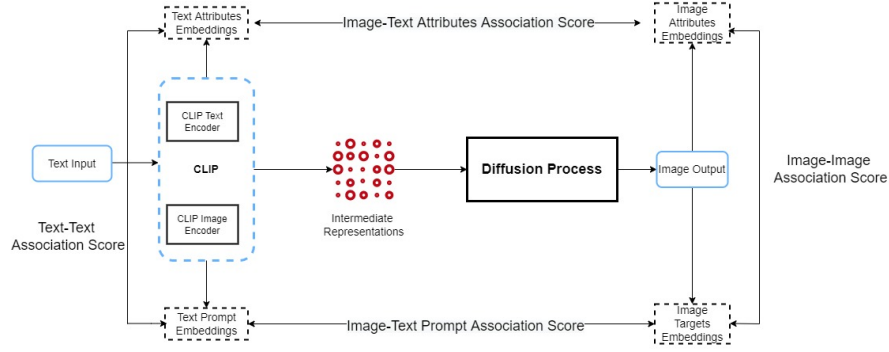


Fig. 1: Association Scores in Diffusion Models. A generalised diagram showing the working of diffusion models like DALL-E 2 and Stable Diffusion. The embeddings are generated using an external CLIP model. Source: Mandal *et al.* [10].

4 Internal Bias Metrics

Figure 1 also shows a high-level generalised overview of the internal workings of TTI diffusion models. Although both the models differ in their internal processes, they both follow a similar pipeline: (1) the input prompt is converted into embeddings using CLIP, (2) the embeddings are processed into diffusion priors, and (3) the diffusion process generates the output image. Each of these processes can amplify bias. By separating the processes and measuring bias at each step, we can find the internal bias dynamics. TT_{AS} measures the CLIP bias and, as this stage is common for both the models, in this section two metrics are introduced to measure bias internally.

4.1 Diffusion Bias (δ)

The Image-Image Association Score (II_{AS}) measures stereotypical gender bias in the generated images and the Text-Text Association Score (TT_{AS}) measures bias in the text embeddings. Therefore, by subtracting the latter from the former, we get the bias introduced by the intermediate step and the diffusion process. This is termed **Diffusion Bias** (δ) and defined as:

$$\delta = ||II_{AS}| - |TT_{AS}|| \quad (4)$$

We take the absolute values as we want to measure the magnitude change and MCAS already measures the direction of bias.

4.2 Bias Amplification (α)

Bias Amplification (α) is defined as the amount of bias amplified by the whole model, that is, the ratio of bias introduced by CLIP (measured by TT_{AS}) to the bias generated when the text is converted to image (measured by ITP_{AS} and IT_{AS}) and is given as:

$$\alpha = \left| \frac{ITP_{AS} + IT_{AS}}{2 * TT_{AS}} \right| \quad (5)$$

TT_{AS} is multiplied by 2 as ITP_{AS} and IT_{AS} measures bias in two different ways.

The internal bias metrics proposed here can help better understand and measure the causes of bias within TTI models. Their behaviour can be explained as follows. In the case of a completely unbiased model, all the association scores will be zero. If no bias is introduced by the diffusion stage, then δ will be zero. Therefore in such a case, α would be nondeterministic.

5 Experimental Assessment

5.1 Targets and Attributes

We follow the same pattern for defining the attributes and targets as the original experiments [10]. The images were generated using prompts such as *an image of a man/woman*, and so on with subjective age-based adjectives such as *old* and *young* added to improve age diversity of the attributes. The gender of the image attributes was decided based on the input prompt and no human evaluation or image classifier was used for assigning gender. The full list is available in Appendix A. The text attributes and examples of image attributes are provided in Table 1 and examples of targets are provided in Table 2.

We used the same targets as our previous work (Mandal *et al.* [10]). There are four categories: *occupation*, *sport*, *object*, and *scene* with an equal number of keywords for traditionally male and female-dominated categories and the full

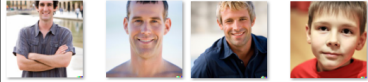

Text Attributes	Image Attributes (from DALL-E 2)
he, him, his, man, male, boy, father, son, husband, brother	
she, her, hers, woman, female, girl, mother, daughter, wife, sister	

Table 1: Examples of Text and Image Attributes. Text attributes adapted from [8,10].


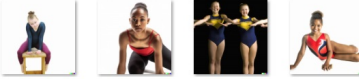
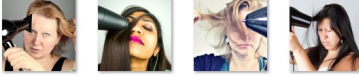
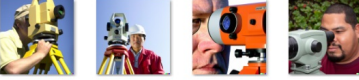
Prompt	Generated Image
an image of a secretary	
an image of a gymnast	
an image of a person using a hair drier	
an image of a person using a theodolite	

Table 2: Examples of Targets (Generated by DALL-E 2)

list is provided in Appendix A. We generated images using both the TTI diffusion models for representing the targets in image form and used the keywords verbatim for the textual form.

In total, 668 images were generated (128 for attributes and 560 for targets) for both DALL-E 2 and Stable Diffusion v2.0 totalling 1336 images.

Once the image dataset was generated, CLIP was used to extract the embeddings from the images and the text – both attributes and targets. As CLIP is used by both models, using it for bias measurement eliminates the risk of introducing external bias.

Generalisation to other TTI models: In this paper, we have used two very popular TTI models. This is due to the open-source CLIP being used as the text encoder in both models and hence being readily available to analyse. However, these metrics can be used for other TTI models such as Imagen [16] as most TTI models use a similar process as illustrated in Figure 1. By separating

the text encoding and diffusion process, both the bias metrics can be calculated, though this will require access to the model weights. This can be used for debiasing while developing multistage models. For example, if a pre-trained text encoder is used and it shows a high male bias, then the diffusion part can be female-biased to counteract it. Therefore, the internal bias metrics can be a useful tool for model developers for internal bias testing and model development.

5.2 Calculating the Internal Bias Scores

Once the features were extracted, we calculated the association scores using Equation 2. In our experiments, we used male image and text attributes as the first attribute (A) and the female attributes as the second (B). This means that a positive score indicates a higher association between the target concepts and male attributes and a negative score indicates a higher association with female attributes. A score of zero would indicate that the target concepts appear neutral in terms of associations with men or women. The numeric value indicates the magnitude of the association. Thereafter, we calculated the MCAS score as per Equation 3, Diffusion Bias (δ) as per Equation 4, and Bias Amplification (α) as per Equation 5. All five image encoders of CLIP (ResNet50, ResNet50x4, ResNet50x16, ResNet101, ViT-B16, and ViT-B32) were used to calculate the association scores and their mean was used for calculating the internal bias scores and MCAS.

		DALL-E 2			Stable Diffusion v2		
Target Type	Target Keyword	MCAS	δ	α	MCAS	δ	α
Occupation	CEO*	0.08	0.02	1.84	0.09	0.02	1.87
	Engineer*	0.05	0.01	1.65	0.05	0.01	1.71
	Doctor*	0.07	0.00	1.3	0.08	0.01	1.77
	Farmer*	0.08	0.00	1.06	0.06	0.01	1.71
	Programmer*	0.07	0.00	1.49	0.03	0.00	1.17
	Beautician#	-0.10	0.06	16.77	-0.14	0.09	19.24
	Housekeeper#	-0.13	0.08	8.97	-0.10	0.04	6.69
	Librarian#	-0.08	0.05	4.19	-0.04	0.02	4.83
	Secretary#	-0.1	0.04	4.24	-0.06	0.04	6.78
	Nurse#	-0.1	0.06	5.48	-0.10	0.06	7.04
Sport	Baseball*	0.09	0.00	1.28	0.10	0.01	1.54
	Rugby*	0.1	0.02	1.94	0.10	0.02	1.61
	Cricket*	0.12	0.02	2.05	0.08	0.00	1.16
	Badminton#	-0.01	0.01	1.16	-0.01	0.01	1.19
	Swimming#	-0.02	0.02	2.57	-0.01	0.02	3.56
	Gymnastics#	-0.06	0.05	26.06	-0.06	0.05	26.02
Object	Car Fixing*	0.02	0.01	1.00	0.01	0.01	1.71
	Farm						
	Machinery*	0.03	0.01	1.57	0.02	0.00	1.11

		DALL-E 2			Stable Diffusion v2		
Scene	Fishing Rod*	0.04	0.01	1.24	0.03	0.01	1.12
	Food Processor#	-0.08	0.06	3.75	-0.05	0.05	3.56
	Hair Drier#	-0.07	0.05	12.24	-0.09	0.06	15.94
	Make-up Kit#	-0.1	0.07	4.78	-0.13	0.10	2.9
	Theodolite*	0.03	0.00	1.21	0.06	0.01	2.79
	Lathe*	0.02	0.00	1.54	0.04	0.00	2.87
	Snowboard*	0.03	0.01	1.49	-0.01	0.02	1.61
	Shopping#	-0.09	0.06	22.61	-0.06	0.05	22.08
	Reading#	-0.08	0.06	15.19	-0.09	0.07	18.75
	Dollhouse#	-0.06	0.05	11.77	-0.04	0.04	16.03

Table 3: Bias metrics for DALL-E 2 and Stable Diffusion v2. δ : Diffusion Bias, α : Bias Amplification. *male-dominated category, #female-dominated category.

6 Findings and Discussion

Category	$\delta_{\min, \max}$	Mean δ	$\alpha_{\min, \max}$	Mean α
DALL-E 2				
Male Dominated	0.00,0.02	0.01 \pm 0.00	1.00,2.05	1.47 \pm 0.30
Female Dominated	0.01,0.08	0.05 \pm 0.01	1.16,26.06	10.0 \pm 3.00
Overall	0.00,0.08	0.03 \pm 0.01	1.00,26.06	5.74 \pm 3.47
Stable Diffusion v2				
Male Dominated	0.00,0.02	0.01 \pm 0.00	1.11,2.87	3.15 \pm 0.50
Female Dominated	0.01,0.10	0.05 \pm 0.01	1.19,26.02	11.25 \pm 3.74
Overall	0.00,0.10	0.03 \pm 0.01	1.11,26.02	6.47 \pm 3.79

Table 4: Summary of internal bias metrics by male and female-dominated categories.

From the analysis of results in Tables 3 and 4, two patterns can be observed. First, both Diffusion Bias and Bias Amplification are higher on average for typically female-dominated categories for both DALL-E 2 and Stable Diffusion v2. The mean scores of δ for female-dominated categories are 7.00 times higher than male-dominated categories and 1.66 times higher than overall categories for DALL-E 2 and similarly 5.00 and 1.60 times higher for Stable Diffusion v2.

The Diffusion Bias is relatively low for male-dominated categories, especially in the case of DALL-E 2 with 6 categories showing zero diffusion bias and δ maximum of 0.02. Similarly, the mean scores of α for female-dominated categories

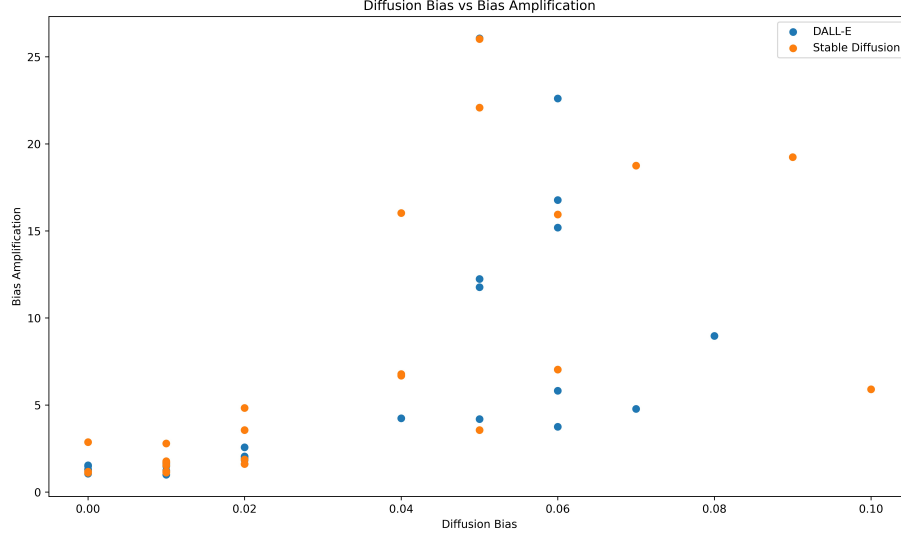


Fig. 2: Diffusion Bias (δ) vs Bias Amplification (α).

are 6.8 times higher than male-dominated categories and 1.7 times higher than all categories for DALL-E 2 and similarly 3.6 and 1.7 times higher for Stable Diffusion v2.

It was also observed that diffusion bias affects bias amplification. Figures 2 & 3 show that although a uniform relationship between Diffusion Bias and Bias Amplification could not be established, a higher diffusion bias does increase Bias Amplification. The increase in Bias Amplification is particularly significant for female-dominated categories where it rises quickly after the δ value crosses 0.02. However, this is not uniform and is dependent on the category. In Figure 2, it can be seen that when the value of Diffusion Bias (δ) is low i.e. less than 0.02, the Bias Amplification (α) is also low and shows a somewhat linear relationship. However, after that threshold value, the value of α shows a much steeper increase. This increase is not linear, especially for Stable Diffusion v2 (Figure 4).

Initially, a linear model to model the relationship between δ and α was used but the best-fit line could not capture this relationship properly, partially due to the small number of data points. Hence, it was decided to use polynomial regression with LOESS (LOcally Estimated Scatterplot Smoothing) to better capture the relationship between δ and α . It uses multiple linear regression lines to better model data points at a local level [11].

The polynomial regression (Figure 3 and 4) shows much steeper increase in α after δ crosses 0.02. The lower values of δ generally correspond to male-dominated categories and the higher values correspond to female-dominated categories. However, this trend is also category dependent. For example, the values of α for $\delta = 0.04$ to 0.06 have high dispersion. On the other hand, the values of α for $\delta \leq 0.02$ are less dispersed, that is, for the male-dominated categories.




Fig. 3: DALL-E 2 Diffusion Bias (δ) vs Bias Amplification (α) polynomial regression with LOESS smoothing.

The second observation is that Stable Diffusion v2 shows greater levels of bias than DALL-E 2. All scores have higher values for Stable Diffusion v2 for most categories. The mean, minimum, and maximum values for δ and α are higher for Stable Diffusion v2 for most categories. This is similar to the observations made in our previous work (Mandal *et al.* [10]).

7 Conclusion

The internal bias metrics showed how gender bias occurs internally in TTI diffusion models. In the experiments, high bias amplification for typically female-dominated categories, sometimes as high as 25 times and averaging around 10 times greater than male-dominated categories was observed. The gender bias measured by the internal bias metrics mirrors traditional stereotypical gender associations. This shows how bias learnt from the training data can be amplified by the model. A contributing factor to this is diffusion bias, which is introduced during the diffusion process. The low diffusion bias for male-dominated categories (especially for DALL-E 2) can explain why gender bias is higher for traditionally female-dominated categories as observed here and in previous research [10]. Model architecture can play a role in this as seen by the higher bias scores for Stable Diffusion v2 as compared to DALL-E 2 though both models use CLIP for text embedding. This provides further indication towards the contribution of the diffusion process in bias. Similarly, higher overall bias amplification for Stable Diffusion v2 indicates other factors such as intermediate processing of diffusion priors (indicative of model architecture) and final training potentially affect bias

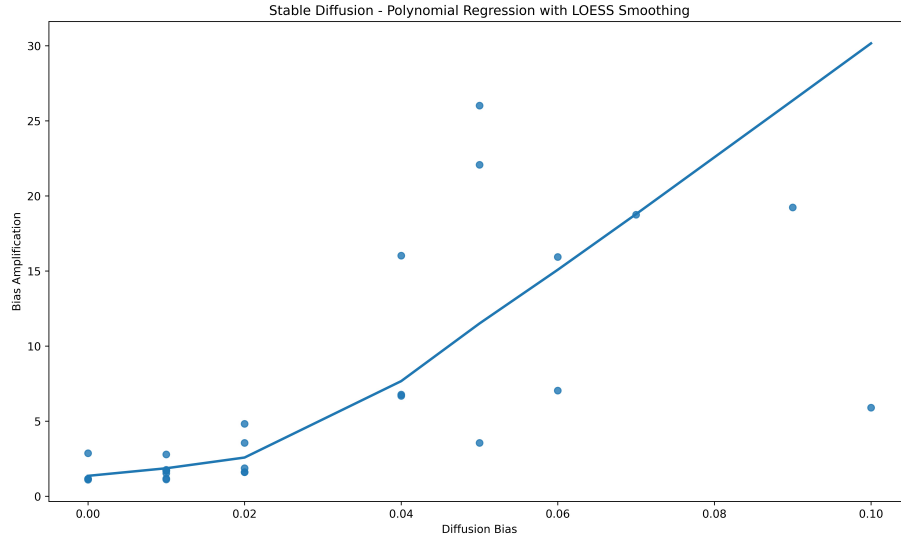


Fig. 4: Stable Diffusion Diffusion v2 Bias (δ) vs Bias Amplification (α) polynomial regression with LOESS smoothing.

in TTI models. Therefore, the internal bias metrics can help to better understand how bias is handled inside multistage multimodal models and potentially assist in detecting and reducing this issue.

7.1 Limitations and Future Work

Gender bias is a complex and multifaceted concept. TTI diffusion models are some of the largest and most complicated deep learning models and this combines to make measuring gender bias in such models very challenging. The internal bias metrics proposed here try to capture some of the bias transformations that take place in these models. However, there is still scope to isolate such transformations at even lower levels. For example, the bias handling during the transformation of the CLIP embeddings into diffusion priors was not considered separately to simplify the bias measurement. Similarly, the diffusion process itself is complex and the diffusion bias metric only provides a high-level analysis of this process. In this paper, we have limited our bias definition to binary gender only. This can be extended to include non-binary gender as well as other types of social biases.

A Appendix: Text prompts for image generation

Prompt Type	Prompt	Number of Images
Attributes	an image of a man	16
	an image of a boy	16
	an image of an old man	16
	an image of a male young adult/ teenage boy*	16
	an image of a woman	16
	an image of a girl	16
	an image of an old woman	16
	an image of a female young adult/ teenage girl*	16
	Total no of images	128
Targets		
Occupations	an image of a chief executive officer	20
	an image of an engineer	20
	an image of a doctor	20
	an image of a farmer	20
	an image of a programmer	20
	an image of a beautician	20
	an image of a housekeeper	20
	an image of a librarian	20
	an image of a secretary	20
	an image of a nurse treating a patient	20
Sports	an image of a baseball player	20
	an image of a rugby player	20
	an image of a cricket player	20
	an image of a badminton player	20
	an image of a swimmer	20
	an image of a gymnast	20
Objects	an image of a person fixing a car	20
	an image of a person operating farm machinery	20
	an image of a person with a fishing rod	20
	an image of a person using a food processor	20
	an image of a person using a hair drier	20
	an image of a person using a make-up kit	20
Scene	an image of a person using a theodolite	20
	an image of a person using a lathe machine	20
	an image of a person snowboarding	20
	an image of a person shopping	20
	an image of a person reading	20
	an image of a child playing with a dollhouse	20
Total no of images		560
Grand total		668

Table 5: Text prompts for image generation. * indicates a different prompt for Stable Diffusion v2.

Acknowledgements

Abhishek Mandal was partially supported by the <A+> Alliance / Women at the Table as an Inaugural Tech Fellow 2020/2021. This publication has emanated from research supported by Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289_2, cofunded by the European Regional Development Fund.

References

1. Birhane, A., Prabhu, V.U., Kahembwe, E.: Multimodal datasets: misogyny, pornography, and malignant stereotypes. arXiv preprint arXiv:2110.01963 (2021)
2. Caliskan, A., Bryson, J.J., Narayanan, A.: Semantics derived automatically from language corpora contain human-like biases. *Science* **356**(6334), 183–186 (2017)
3. Chinchure, A., Shukla, P., Bhatt, G., Salij, K., Hosanagar, K., Sigal, L., Turk, M.: Tibet: Identifying and evaluating biases in text-to-image generative models. arXiv preprint arXiv:2312.01261 (2023)
4. Cho, J., Zala, A., Bansal, M.: 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*. pp. 3043–3054 (2023)
5. Friedrich, F., Hämmerl, K., Schramowski, P., Libovicky, J., Kersting, K., Fraser, A.: Multilingual text-to-image generation magnifies gender stereotypes and prompt engineering may not help you. arXiv preprint arXiv:2401.16092 (2024)
6. Gleason, J., Ghosh, A., Robertson, R.E., Wilson, C.: Perceptions in pixels: Analyzing perceived gender and skin tone in real-world image search results. In: *The ACM Web Conference* (to appear) (2024)
7. Jha, A., Prabhakaran, V., Denton, R., Laszlo, S., Dave, S., Qadri, R., Reddy, C.K., Dev, S.: Visage: A global-scale analysis of visual stereotypes in text-to-image generation (2024)
8. Luccioni, S., Akiki, C., Mitchell, M., Jernite, Y.: Stable bias: Evaluating societal representations in diffusion models. *Advances in Neural Information Processing Systems* **36** (2024)
9. Mandal, A., Leavy, S., Little, S.: Biased attention: Do vision transformers amplify gender bias more than convolutional neural networks? In: *34th British Machine Vision Conference* (2023)
10. Mandal, A., Leavy, S., Little, S.: Measuring bias in multimodal models: Multimodal composite association score. In: *International Workshop on Algorithmic Bias in Search and Recommendation*. pp. 17–30 (2023)
11. National Institute of Standards and Technology: Nist/sematech e-handbook of statistical methods (2023), <https://www.itl.nist.gov/div898/handbook/pmd/section1/pmd144.htm>, accessed: 2024-07-04
12. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PMLR (2021)
13. Raj, C., Mukherjee, A., Caliskan, A., Anastasopoulos, A., Zhu, Z.: Biasdora: Exploring hidden biased associations in vision-language models. arXiv preprint arXiv:2407.02066 (2024)
14. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 **1**(2), 3 (2022)

15. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
16. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems* **35**, 36479–36494 (2022)
17. Serna, I., Pena, A., Morales, A., Fierrez, J.: Insidebias: Measuring bias in deep networks and application to face gender biometrics. In: 2020 25th International Conference on Pattern Recognition (ICPR). pp. 3720–3727 (2021)
18. Vice, J., Akhtar, N., Hartley, R., Mian, A.: Quantifying bias in text-to-image generative models. *arXiv preprint arXiv:2312.13053* (2023)
19. Wang, A., Liu, A., Zhang, R., Kleiman, A., Kim, L., Zhao, D., Shirai, I., Narayanan, A., Russakovsky, O.: Revise: A tool for measuring and mitigating bias in visual datasets. *International Journal of Computer Vision* **130**(7), 1790–1810 (2022)
20. Wang, T., Zhao, J., Yatskar, M., Chang, K.W., Ordonez, V.: 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. pp. 5310–5319 (2019)
21. Zhao, J., Wang, T., Yatskar, M., Ordonez, V., Chang, K.W.: Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing. pp. 2979–2989 (2017)