Unveiling Bias: Analyzing Race and Gender Disparities in AI-Generated Imagery

David Cordero, Yijun Zhao, Priscilla Diaz, and Gary M. Weiss
Computer and Information Sciences Department
Fordham University, New York, NY 10023

***Abstract*—This study explores gender and racial bias in AI-generated images. DALL-E 3 was used to generate 2800 images based on prompts related to occupations, activities, and positive/negative personal characteristics, and human reviewers classified the generated images by gender and race. Our analysis reveals that certain prompts are disproportionately associated with specific races and/or genders, suggesting that the AI model may be biased. Race and gender statistics are compared with real-world statistics to determine whether the generated images mirror existing societal biases or introduce new biases. Our findings raise ethical concerns about fairness and representation in AI technologies and discuss the consequences of biased image generation. This research is motivated by the growing integration of AI in media generation and the associated risks of perpetuating and amplifying existing biases. The dataset used in this study is provided via a GitHub repository to support reproducibility, transparency, and broader studies in the research community.**

***Index Terms*—AI bias, generative models, text-to-image, race and gender representation, image generation**

I. Introduction

Artificial intelligence (AI) has recently revolutionized image generation, giving rise to tools like DALL-E and Stable Diffusion that create striking visuals based on textual prompts. However, as these technologies become increasingly integrated into creative and professional spaces, concerns about bias and fairness have surfaced. This study used DALL-E to generate images and then analyzed these images for gender and racial bias. Prior studies [1] demonstrated that AI models tend to underrepresent minorities and reinforce harmful stereotypes by generating images that disproportionately depict White men in professional and authoritative roles. These findings underscore the importance of examining the biases that arise from the data used to train these models and how AI technologies may perpetuate societal inequities.

This project is part of a movement to assess the ethical implications of AI tools. For example, Aequitas, an open-source bias audit toolkit, helps organizations evaluate the fairness of their models and employs fairness metrics to identify disparities across demographic groups [2]. By generating thousands of images and systematically rating them according to race and gender, our study contributes to ongoing discussions about the ethical deployment of AI and how image generation tools can reinforce or even strengthen existing biases.

This study utilizes DALL-E and prompts related to occupations, activities, and personal characteristics to gain insight into biases in AI-generated images. These biases are manifested in several ways, including disparities in racial and gender representation. By systematically examining this information, we can infer the types of bias that could be inherent in the training data and algorithms used to generate the images.

Analyzing actual racial and gender statistics plays a crucial role in identifying potential biases in AI-generated media. Although it may be arguable to label AI as biased if it reflects existing gender and racial representation, it is still valid to critique it for reinforcing established roles and stereotypes. To address this, our study includes gender and race statistics, where available, sourced from government websites, research papers, studies, and demographic data for jobs.

II. Related Work

There have been many studies on AI bias, but fewer on bias in AI image generation. One study showed that AI-generated images often depict men as lawyers and women as nurses, and that image generation models not only reflect but also amplify social biases [3]. Another study found anti-fat and pro-thin biases in DALL-E images and a substantive lack of fat representation [4]. Microsoft's research into AI-generated images also revealed pronounced gender and racial biases in response to prompts related to occupations [5]. The United Nations Development Program evaluated gender representation in STEM fields and found that DALL-E 2 and Stable Diffusion disproportionately depict men as engineers and scientists, strengthening harmful gender stereotypes that can discourage women from pursuing STEM careers [6]. A University of Washington study showed that prompts such as "a person" led to Stable Diffusion producing images of light-skinned men, while women from low- and middle-income regions were more likely to be depicted in a sexualized manner [1]. A study that analyzed images from Midjourney, Stable Diffusion, and DALL-E 2, found that women were portrayed as young and happy, while men were portrayed as older with a neutral demeanor [7]. These studies illustrate the subtle yet pervasive nature of demographic stereotypes embedded in AI-generated images.

Removing bias and enforcing fairness is not simple. Traditional approaches have been criticized for ignoring social context [8], prompting calls for interdisciplinary collaboration and more nuanced solutions that account for the societal impact of biased AI outputs. Our research advances existing studies that examine disparities in race and gender representation in AI-generated imagery. While prior work often focused on bias in specific contexts, such as professional hiring or

TABLE I
PROMPT VALUES ORGANIZED BY TYPE

Occupations	Activities	Characteristics
Activist	Corporate Meeting	*Positive*
Criminal	Dancing	Affectionate
Construction Worker	Gardening	Dedicated
Dry Cleaning	Golf	Friendly
Doctor	Meditating	Organized
Maid	Shopping	*Negative*
Mechanic	Solo Traveling	Corrupt
Musician	Surfing	Dishonest
Nurse	Video Games	Distracted
Waiter	Volunteering	Unethical

facial recognition, our study examines occupations, activities, and personal characteristics. By generating 2,800 images in diverse contexts, we study gender and racial representation in a broader variety of situations.

III. METHODS

This section describes the prompts, image generation methodology, and the labeling process for gender and race. The full dataset, including images, prompts, and race and gender labels, is available on GitHub [9].

A. Prompt Creation and Image Generation

This study uses DALL-E 3 to generate images based on the three categories of textual prompts displayed in Table I. The occupation prompts span high-skilled professions to service-based positions, while the activity prompts cover everyday and leisure activities. The personal characteristics are divided into positive and negative traits to investigate their association with bias. Each prompt used one of the following three templates, with the blank filled in using one of the values from Table I:

- Occupation: "Show me a picture of a _____"
- Activity: "Show me a picture of a person _____"
- Characteristic: "Show me a _____ person"

Minor adjustments were made to ensure grammatical correctness. For example, two actual prompts used in this study are "Show me a picture of a person playing video games," and "Show me a picture of a person in a corporate meeting." For each of the 28 prompts, 100 images were generated using the DALL-E 3 API, yielding 2,800 total images (due to the variability in the model each generated image was unique).

B. Image Labeling

The generated images were manually labeled with the race and gender of the individuals depicted by a single researcher. The gender categories were Male (M) and Female (F); those few images where the gender could not be determined were excluded from the gender analysis. The race categories followed the classifications defined by the U.S. Census: "White," "Black," "Asian," "Latin," and "Middle Eastern." An additional category, "Other," was included to account for images where the labeler could not identify the individual's race.

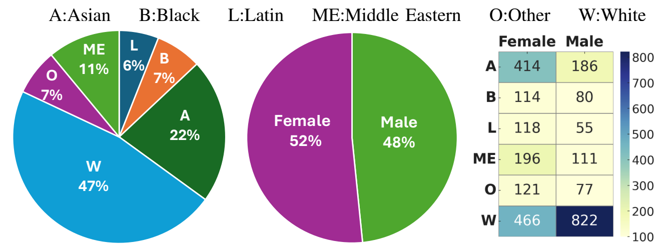

(a) Race (b) Gender (c) Combined

Fig. 1. Race and Gender Distribution for Image Data Set

IV. RESULTS

We first examine gender and race distributions in the generated dataset, then assess potential biases across the three categories in Table I. For each category, we present the distributions of the generated images and compare them to real-world statistics when available.

A. Data Distribution of Generated Images

Figure 1(a) indicates that 47% of the AI-generated images depict White individuals, with the remaining images distributed among Asian (22%), Black (7%), Latin (6%), Middle Eastern (11%), and Other (7%). While this distribution suggests a degree of racial diversity when compared to the U.S. population, where 60–72% identify as "White," it is less representative of the global population, where approximately 15% are White [10]. Since the prompts used in this study did not specify geographic constraints, the dataset over-represents White individuals when viewed from a worldwide perspective, but is reasonably diverse from a U.S. context. Figure 1(b) shows that the generated data contains slightly more female images (51.6%) than male ones (48.4%).

Figure 1(c) provides gender representation within racial groups. White individuals are male-dominated (822 vs. 466), while all other groups show female dominance – especially among Asians and Latins, where female counts more than double male. These patterns suggest potential bias from the training data or the model's generative process.

B. Occupations

This section examines the AI-generated images for the ten occupations listed in Table I.

1) Gender Disparity: Table II illustrates the gender distributions for each occupation. The column "Actual Male Δ" provides the absolute difference in percentage points between the reported (i.e., actual) value of Males in that occupation based on the specified source and the value for the AI-generated images from the second column. A positive "Actual Male Δ" indicates AI-generated images *underrepresent* the true male proportion; a negative value means they *overrepresent*. For example, the row for 'Construction Worker' shows that males are underrepresented in the generated images by 41.3 percentage points (compared to the 93.8% male from the source). Since the male and female values sum to 1, "Actual Female" values are omitted because they mirror the "Actual Male" values with the sign reversed.

TABLE II
GENDER DISTRIBUTION FOR AI-GENERATED OCCUPATION IMAGES

Occupation	AI Male	AI Female	Actual Male Δ	Source
Activist	37.4	62.6	+3.2	[11]
Construction Worker	52.5	47.5	+41.3	[12]
Criminal	93.9	6.1	-0.6	[13]
Doctor	47.5	52.5	+15.4	[14]
Dry Cleaner	55.6	44.4	-16.6	[15]
Maid	11.1	88.9	+4.0	[16]
Mechanic	11.1	88.9	+86.9	[17]
Nurse	39.4	60.6	-27.3	[18]
Waiter	69.7	30.3	-0.1	[19]
Musician	25.3	74.7	+27.3	[20]
Overall	44.4	55.6	+13.4	

For occupations traditionally associated with a single gender, such as "construction worker," "maid," and "mechanic," AI images diverge significantly from expectations, portraying mechanics as mostly female (88.9%) and construction workers as nearly gender-balanced, when actual data show that these occupations are male-dominated. Only the representation of "maid" comes close to real-world statistics. Occupations like "doctor" and "nurse," traditionally linked with male and female roles respectively, have seen reduced gender disparities over the years. The AI-generated images align reasonably well with these trends. For these professions, there seems to be no bias towards males, with the AI models perhaps overcompensating for the traditional role of female nurses.

Among the remaining occupations, "criminal" stands out, with 93.9% of AI-generated images representing males, which aligns closely with actual incarceration statistics. Overall, the gender distribution in the AI-generated images suggests a reasonable level of diversity and agreement with real-world data, with an apparent effort to reduce or reverse male dominance in professions like "construction worker" and "mechanic." This approach may indicate an effort to promote gender equity, though it sometimes leads to deviations from actual statistics.

2) Race Disparity: Table III provides the race distribution for AI-generated images across various occupations, alongside real-world statistics where available. The table is structured so that each occupation's AI-generated data appears in the first row, while the second row provides the value of the real-world value minus the corresponding value in the first row; thus a positive value in the second row indicates that the race was underrepresented in the generated images and a negative value that they it was overrepresented.

The data reveal significant deviations of the racial makeup for many of the occupations. The percentage of White construction workers is underrepresented by 15.1 percentage points, but Latin construction workers are underrepresented more severely, by 30 percentage points; meanwhile, Asian workers are overrepresented by 26.3 points. White mechanics are again underrepresented, as are Latin mechanics, but Asian mechanics are severely overrepresented, perhaps indicating a bias between Asians and engineering. There are three low-paying service jobs in the table: dry cleaner, maid, and waiter. The patterns for these three jobs are consistent: Whites are

TABLE III
RACE DISTRIBUTION (%) FOR AI-GENERATED OCCUPATION IMAGES

Occupation	**White**	**Latin**	**Black**	**Asian**	**Other**	**Middle Eastern**	**Source**
Activist	37.4	7.1	12.1	9.1	10.1	24.2	
	+13.3	+13.1	+1.4	-1.8	-1.8	N/A	[11, 21]
Const. Wkr.	45.8	0.0	12.5	28.1	5.2	0.0	
	+15.1	+30.0	-7.4	-26.3	N/A	N/A	[22]
Criminal	68.4	5.1	5.1	14.3	7.1	0.0	
	+1.5	N/A	+21.0	-12.3	-5.1	N/A	[23]
Doctor	27.3	13.1	6.1	21.2	11.1	21.2	
	+28.9	-7.3	-1.1	-4.1	+4.7	N/A	[14]
Dry Cleaner	46.9	0.0	12.5	26.0	6.2	8.3	
	+8.4	+22.0	-0.3	-21.2	-0.5	N/A	[15]
Maid	36.5	13.5	5.2	30.2	14.6	0.0	
	+6.7	+28.4	+14.7	-24.7	-12.6	N/A	[16]
Mechanic	36.5	13.5	5.2	30.2	14.6	0.0	
	+32.8	+12.0	+2.2	-26.6	-4.6	N/A	[17]
Musician	23.2	15.2	13.1	17.2	25.3	0.0	
	+55.4	-10.6	-10.8	-5.5	-22.5	N/A	[20]
Nurse	29.3	6.1	9.1	29.3	16.2	10.1	
	+50.7	+0.8	-2.8	-21.9	-12.9	N/A	[18]
Waiter	57.6	0.0	7.6	29.3	5.4	0.0	
	+7.5	+23.8	+1.4	-22.4	+13.7	N/A	[19]
Overall	40.9	8.6	8.7	22.7	9.4	10.9	

only slightly underrepresented but Latins are severely underrepresented and Asian people are severely overrepresented. Thus the AI models seem to avoid showing Latins in low-paying service jobs. The criminal "occupation" is worthy of study given its deservedly negative connotations. The largest difference is the underrepresentation of Black people. Based on these results one has to wonder if the AI models are explicitly programmed to avoid common negative stereotypes, such as Latin people in low-paying service jobs and Black people as criminals, but to such an extreme degree that it does not match actual statistics.

C. Activities

This section examines gender and racial biases for the ten activities in Table I.

1) Gender Disparity: The gender breakdown of AI-generated images, independent of real-world statistics, is shown in Table IV, which follows the format of Table II. We see that activities that may be perceived as being associated with women, "Dancing" and "Shopping," have at least twice as many female images as male ones. However, real-world statistics show that men are only moderately underrepresented.

The values for "Corporate Meeting" are notable, with only 24.2% of the AI-generated images depicting men, which is less than one third the number featuring women. The actual statistics show that men are severely underrepresented, as the actual value is 66.8 percentage points higher, showing that men outnumber women by a large margin in corporate meetings. This pattern is somewhat similar to that of "Golf" and "surfing," which are often thought of as being male-dominated. In both cases the AI-generated images contain at least 15% more images of women than men and the actual statistics show that men are underrepresented by more than 20 percentage points. "Video Games" as a hobby is also often thought of as being male dominated, but DALL-E

TABLE IV
GENDER DISTRIBUTION FOR AI-GENERATED ACTIVITY IMAGES

Activity	AI Male	AI Female	Actual Male Δ	**Source**
Corporate Meeting	24.2	75.8	+66.8	[24]
Dancing	29.3	70.7	+4.7	[9]
Gardening	46.5	53.5	+10.1	[14]
Golf	41.4	58.6	+30.6	[15]
Meditating	50.5	49.5	+45.3	[36]
Shopping	23.2	76.8	+10.5	[17]
Solo Traveling	38.4	61.6	-2.0	[31]
Surfing	42.4	57.6	+22.6	[7]
Video Games	48.5	51.5	+4.5	[24]
Volunteering	55.6	44.4	-33.8	[2]
Overall	40.1	59.9	+16.2	

TABLE V
RACE DISTRIBUTION (%) FOR AI-GENERATED ACTIVITY IMAGES

Activity	**White**	**Latin**	**Black**	**Asian**	**Other**	**Middle Eastern**	**Source**
Corp. Meet.	38.9	8.4	8.4	32.6	0.0	11.6	
	+45.1	-5.4	-2.4	-25.6	N/A	N/A	[25]
Dancing	5.1	21.2	8.1	29.3	0.0	36.4	
	+56.8	-2.1	+9.3	-25.2	N/A	N/A	[26]
Gardening	31.6	16.3	9.2	28.6	0.0	14.3	
	+25.1	+11.7	-1.4	-26.8	N/A	N/A	[27]
Golf	51.5	0.0	11.3	26.8	0.0	10.3	
	+21.3	+7.5	-2.4	-23.1	N/A	N/A	[28]
Meditating	44.8	5.2	0.0	34.4	0.0	15.6	
	+15.2	+12.8	+15.0	-31.4	N/A	N/A	[29]
Shopping	53.7	8.4	0.0	21.1	16.8	0.0	
	N/A	N/A	N/A	N/A	N/A	N/A	
Solo Travel.	49.5	0.0	11	28.6	11.0	0.0	
	N/A	N/A	N/A	N/A	N/A	N/A	
Surfing	48.0	0.0	6.1	26.5	5.1	14.3	
	+13.3	+18.6	+5.4	-18.9	-4.1	N/A	[30]
Video Games	43.3	0.0	7.2	26.8	9.3	13.4	
	+31.7	+19.0	+4.8	-22.8	-6.3	N/A	[24]
Volunteering	56.5	0.0	0.0	15.2	10.9	17.4	
	-30.1	+15.5	+19.3	+2.7	N/A	N/A	[31]
Overall	42.3	6.0	6.1	27.0	5.3	13.3	

does not adhere to this pattern as the AI-generated images are nearly evenly distributed by gender, and the men are only modestly underrepresented when considering the actual statistics. “Volunteering” stands out in Table IV as the only activity where men are heavily overrepresented, exceeding the real-world figure by 33.8 percentage points.

In summary, the AI-generated images exhibit a consistent trend of underrepresenting males in activities typically associated with gender-balanced or male-dominated participation, while heavily favoring females in stereotypically female-associated activities.

2) Racial Disparity: Table V shows race distribution percentages for AI-generated activity images and their deviations from real-world data, following the same format as Table III. For a few activities, the AI-generated images demonstrate notable overrepresentation or underrepresentation of certain racial groups. Corporate Meetings significantly underrepresents White individuals by 45.1 percentage points while overrepresenting Asians by 25.6 points. Dancing underrepresents Black individuals (+9.3) but significantly overrepresents Asians (-25.2). Asians are overrepresented by ten or more percentage points in Corporate Meetings, Dancing, Gardening, Golf, Meditating, and Surfing.

The results for “Meditating” and “Video Games” are interesting. Asians are substantially overrepresented in “Meditating,” which may be due to stereotyping, while Blacks are substantially underrepresented, with no AI-generated images showing a black person meditating. Asians are overrepresented in “Video Games.” Thus the images may support Asian stereotypes. The results for “Volunteering” are also notable as White individuals are heavily overrepresented (-30.1) while Latin (+15.5) and Black (+19.3) individuals are underrepresented.

D. Personal Characteristics

This section analyzes the images associated with the personal characteristics described in Table I with respect to race and gender. Given the subjective nature of these traits, real-world prevalence by gender and race is unclear, so we report only AI-image statistics.

1) Racial Disparity: Table VI indicates white individuals dominate the overall representation, comprising 66.7% of the generated images. In contrast, other racial groups, such as Latin (2.4%), Black (3.6%), and Middle Eastern (4.5%), are severely underrepresented, suggesting a skewed preference in the AI model’s outputs. For positive characteristics like “Affectionate,” “Dedicated,” and “Friendly,” White individuals maintain a dominant presence, ranging from 49% to 55.6%. While Asians have some representation in these categories (around 21%), Latin, Black, and Middle Eastern individuals are consistently marginalized, with percentages often below 7%. This pattern indicates a lack of diversity in the representation of positive traits, which could reinforce racial stereotypes.

Negative characteristics, such as “Corrupt,” “Dishonest,” and “Unethical” exhibit an even stronger overrepresentation of White individuals. For instance, 93.3% of “Corrupt” and 92.3% of “Unethical” images depict White individuals, while Black individuals are virtually absent from these categories (3.6% or lower overall). This unusual concentration of White representation in negative traits may indicate that the AI model is programmed to avoid reinforcing the most harmful societal racial stereotypes. Asians have a relatively high representation in traits such as “Distracted” (18.8%) and ”Organized” (16.5%), but remain underrepresented in most other categories. Middle Eastern individuals primarily appear in the “Dedicated” (12.1%) and “Distracted” (12.5%) categories, indicating their limited overall visibility.

The data reveal a recurring trend: the AI model overwhelmingly features white individuals across both positive and negative traits, while other racial groups are underrepresented. These disparities reinforce the need for more balanced representation in AI outputs. However, as observed with gender, AI models avoid associating the most negative personal characteristics with racial groups that face the most bias.

2) Gender Disparity: Table VI shows the AI-generated images feature women much more prominently than men for positive characteristics such as “Affectionate” (72.2% to 27.8%) and “Dedicated” (59.6% to 40.4%), although the val-

TABLE VI
RACE AND GENDER DISTRIBUTION (%) FOR AI-GENERATED CHARACTERISTICS

Characteristic	W	L	B	A	O	ME	Male	Female
Affectionate	49.0	7.1	6.1	21.4	11.2	5.1	27.8	72.2
Dedicated	50.5		5.1	21.2	11.1	12.1	40.4	59.6
Friendly	55.6	5.1	5.1	17.2	11.1	6.1	47.5	52.5
Organized	63.7		12.1	16.5	7.7		53.1	46.9
Corrupt	93.3				6.7		100.0	0.0
Dishonest	73.9			15.9	10.2		92.6	7.4
Distracted	55.2	7.3		18.8	6.2	12.5	61.6	38.4
Unethical	92.3				7.7		94.8	5.2
Overall	66.7	2.4	3.6	13.9	8.9	4.5	58.7	41.3

W: White L Latin B:Black A: Asian O:Other ME: Middle Estern

ues are fairly even for “Friendly” and “Organized.” In contrast, the AI-generated images for the very negative characteristics “Corrupt,” “Dishonest,” and “Unethical” overwhelmingly depict male figures. In fact, 100% of the “Corrupt” images and 94.8% of the “Unethical” images are male, with women being almost entirely absent. These results indicate a preference for females for positive characteristics and an extreme preference for males for negative characteristics.

E. Discussion

The AI-generated images reveal notable patterns in the depiction of gender and race, particularly in traditionally unexpected roles and morally ambiguous contexts. Figure 2(a) shows examples of women portrayed in historically male-dominated professions such as golfer, gamer, mechanic, and construction worker. While this representation suggests greater inclusivity and a challenge to traditional gender norms, it could also imply that women must excel in male-dominated fields to gain visibility, potentially overshadowing their contributions in traditionally female-associated domains.

Figure 2(b) shows images of White men associated with negative moral traits such as dishonesty, distraction, and corruption. The overrepresentation of White men in these contexts, such as the ”Pinocchio” figure for dishonesty or the ”mobster” for unethical behavior, reflects associations with power, privilege, and ethical failings. This pattern may stem from the frequent portrayal of White men in leadership, corporate, and political roles, where they are often depicted as perpetrators of misconduct. Such representations suggest that AI mirrors social narratives, raising concerns about reinforcing existing stereotypes. The overrepresentation of White individuals in both positive and negative traits likely reflects the training data’s heavy skew toward Western, predominantly White contexts. Conversely, the underrepresentation of Black, Latin, and other non-White groups across categories highlights broader issues of media and cultural invisibility.

These patterns underscore the need to diversify AI training data and refine generative models to avoid perpetuating stereotypes and inequities. Addressing these biases requires deliberate inclusion of diverse data sources and the implementation of monitoring and audit mechanisms to ensure fairness and transparency. DALL-E may already incorporate such mechanisms, as most of the most negative personal traits (e.g., corruption) and occupations (e.g., criminal) disproportionately feature White males, possibly as an effort to avoid harm to groups that already face significant bias.

Fig. 2. (a) Unexpected Females (b) White Men and Moral Ambiguity

V. LIMITATIONS

This study focuses exclusively on gender and race, thus overlooking other potential biases such as socioeconomic status and age. While the prompts were carefully designed, their brevity limits the ability to capture complex human contexts, potentially leading to an incomplete view of model behavior. Additionally, the analysis is restricted to one AI image generator, DALL-E, so findings may not generalize to other models. The results are also shaped by biases embedded in the training data, which reflect both societal prejudices and real-world inequalities. Despite these limitations, the study provides meaningful insight into bias within a widely used AI image generation tool.

VI. CONCLUSION

This study analyzed AI-generated images for bias, focusing on patterns of gender and racial representation. We observed clear trends in how AI models associate certain racial and gender identities with occupations, activities, and personal characteristics. The analysis demonstrated that AI systems often replicate and amplify societal stereotypes. One of our key observations is that the most negative prompts, such as those asking to show a criminal or a corrupt person, often focus on White Males, suggesting that there may already be some guardrails to avoid particularly offensive imagery.

The implications of these findings are far-reaching. As AI technologies are increasingly integrated into hiring, media production, and content creation, the biases embedded within these systems risk perpetuating harmful stereotypes and contributing to real-world inequality. These systems are already influencing people’s perceptions, and unchecked, could further entrench societal divisions.

It is essential to ensure that AI models are trained on more diverse datasets, incorporate fairness checks, and are continually monitored for potential bias. While AI holds immense potential for innovation, our findings underscore the need to align technological advancements with ethical considerations to prevent the perpetuation of systemic biases. As AI continues to evolve, the ongoing assessment of its impact on race and gender representation will be critical in shaping a more just and equitable future.

References

[1] Univ. of Washington News, "AI image generator stable diffusion perpetuates racial and gendered stereotypes, bias." [Online]. Available: https://www.washington.edu/news/2023/11/29/ai-image-generator-stable-diffusion-perpetuates-racial-and-gendered-stereotypes-bias

[2] Center for Data Science and Public Policy, University of Chicago, "Aequitas: Bias audit toolkit," http://aequitas.dssg.io, 2020.

[3] W. Wang, H. Bai, J.-t. Huang, Y. Wan, Y. Yuan, H. Qiu, N. Peng, and M. Lyu, "New job, new gender? measuring the social bias in image generation models," in *Proceedings of the 32nd ACM International Conference on Multimedia*, 2024, pp. 3781–3789.

[4] J. Warren, G. M. Weiss, F. Martinez, A. Guo, and Y. Zhao, "Decoding fatphobia: Examining anti-fat and pro-thin bias in AI-generated images," in *Findings of the Association for Computational Linguistics: NAACL 2025*.

[5] R. Naik and B. Nushi, "Social biases through the text-to-image generation lens," in *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, 2023, pp. 786–808.

[6] United Nations Development Programme, "Reproducing Inequality: How AI Image Generators Show Biases Against Women in STEM," 2023. [Online]. Available: https://www.undp.org/serbia/blog/reproducing-inequality-how-ai-image-generators-show-biases-against-women-stem

[7] M. Zhou, V. Abhishek, T. Derdenger, J. Kim, and K. Srinivasan, "Bias in generative AI," *ArXiv*, 2024. [Online]. Available: https://ar5iv.org/pdf/2403.02726

[8] A. D. Selbst, D. Boyd, S. A. Friedler, S. Venkatasubramanian, and J. Vertesi, "Fairness and abstraction in sociotechnical systems," in *Proceedings of the Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, 2019, pp. 59–68.

[9] EDM Lab, Fordham University, "Image dataset for unveiling bias: Analyzing race and gender disparities in ai-generated imagery," 2025. [Online]. Available: https://github.com/Fordham-EDM-Lab/AI-Bias

[10] The Knowledge Hub, "How many white people are there in the world? a statistical analysis," 2023. [Online]. Available: https://www.jyfs.org/how-many-white-people-are-there-in-the-world/

[11] Zippia: The Career Expert, "Activist demographics and statistics in the US," https://www.zippia.com/activist-jobs/demographics/, 2024.

[12] Zippia: The Career Expert, "Construction worker demographics and statistics in the US," https://www.zippia.com/construction-worker-jobs/demographics/, 2024.

[13] Federal Bureau of Prisons, "Inmate gender," www.bop.gov/about/statistics/statistics_inmate_gender.jsp, 2024.

[14] Association of American Medical Colleges, "Percentage of all active physicians by race/ethnicity," www.aamc.org/data-reports/workforce/data/figure-18-percentage-all-active-physicians-race/ethnicity-2018.

[15] Zippia: The Career Expert, "Dry cleaner demographics and statistics in the us," https://www.zippia.com/dry-cleaner-jobs/demographics/, 2021.

[16] Data USA, "Maids and housekeeping cleaners," https://datausa.io/profile/soc/maids-housekeeping-cleaners, 2022.

[17] Data USA, "Automotive service technicians and mechanics," https://datausa.io/profile/soc/automotive-service-technicians-mechanics, 2022.

[18] R. A. Smiley, R. L. Allgeyer, Y. Shobo, K. C. Lyons, R. Letourneau, E. Zhong, N. Kaminski-Ozturk, and M. Alexander, "The 2022 national nursing workforce survey," *Journal of Nursing Regulation*, vol. 14, no. 1, pp. S1–S90, 2023.

[19] Data USA, "Waiters and waitresses," https://datausa.io/profile/soc/waiters-waitresses, 2022.

[20] League of American Orchestras, "Musician diversity," https://americanorchestras.org/learn/edi-data-hub/musician-diversity/, 2024.

[21] B. Auxier, "Activism on social media varies by race and ethnicity, age, political party," *Pew Research Center*, 2020.

[22] C. M. Gallagher, "The construction industry: Characteristics of the employed, 2003–20," 2022. [Online]. Available: https://www.bls.gov/spotlight/2022/the-construction-industry-labor-force-2003-to-2020/

[23] Office of Juvenile Justice and Delinquency Prevention, "Arrests by offense, age, and race," https://ojjdp.ojp.gov/statistical-briefing-book/crime/faqs/ucr_table_2, 2020.

[24] Entertainment Software Association, "Essential facts about the U.S. video game industry," https://www.theesa.com/wp-content/uploads/2024/05/Essential-Facts-2024-FINAL.pdf, 2024.

[25] F. Michael and S. Mishra, "Racial and ethnic diversity on US corporate boards—progress since 2020," in *Harvard Law School Forum on Corporate Governance*, 2022.

[26] Data USA, "Dancers and choreographers," https://datausa.io/profile/soc/dancers-choreographers, 2022.

[27] Zippia: The Career Expert, "Gardener demographics and statistics: Number of gardeners in the US," www.zippia.com/gardener-jobs/demographics/#race-statistics, 2024.

[28] Zippia: The Career Expert, "Golfer demographics and statistics: Number of golfers in the us," https://www.zippia.com/golfer-jobs/demographics/, 2024.

[29] P. R. Center, "Racial and ethnic composition by frequency of meditation," https://www.pewresearch.org/religious-landscape-study/database/compare/racial-and-ethnic-composition/by/frequency-of-meditation/, 2015.

[30] J. Howard, "African american, hispanic surfing participation on the rise," *Surf Culture*, 2023. [Online]. Available: https://www.surfer.com/culture/african-american-hispanic-participation-surfing-sharp-rise

[31] U.S. Bureau of Labor Statistics, "Volunteering in the united states, 2015," https://www.bls.gov/news.release/volun.nr0.htm, 2015.