

# Bias in Generative AI (Work in Progress)

**Mi Zhou**

UBC Sauder School of Business  
[mi.zhou@sauder.ubc.ca](mailto:mi.zhou@sauder.ubc.ca)

**Vibhanshu Abhishek**

UCI Paul Merage School of Business  
[vibs@uci.edu](mailto:vibs@uci.edu)

**Kannan Srinivasan**

CMU Tepper School of Business  
[kannans@cmu.edu](mailto:kannans@cmu.edu)

## Abstract

*This study analyzed images generated by three popular generative artificial intelligence (AI) tools – Midjourney, Stable Diffusion, and DALL·E 2 – representing various occupations to investigate potential biases in AI generators. Our current analysis revealed pervasive gender and racial biases across all three AI tools. First, all three AI generators exhibited bias against women. The average percentage of women in portraits of occupations created by the three AI generators was 23%, 35%, and 42%, respectively, all significantly lower than that of men. Second, all three AI generators were biased against Black people. The average percentage of Black individuals in images created by these three generators was only 9%, 5%, and 2%, respectively, all significantly lower than that of White individuals. Third, the bias in generative AI spanned all job zones, with bright-outlook occupations and STEM occupations facing relatively less bias. Fourth, all three AI generators displayed gender stereotypes in facial expressions and appearances; for example, women were shown with more smiles and happiness, while men were depicted with more neutral expressions and anger. Although the extent of bias varied depending on the model, the direction of bias remained consistent in both commercial and open-source AI generators. As these tools become commonplace, our study highlights the urgency to identify and mitigate bias in generative AI, reinforcing the commitment to ensuring that AI technologies benefit all of humanity for a more inclusive future.*

*Keywords: Generative AI, Text-To-Image, Bias, Equality, Diversity, and Inclusion*

## 1. Introduction

Artificial intelligence (AI) is an indisputable catalyst of societal transformation, reshaping both our professional and personal landscapes. A pivotal player in this transformation is generative AI, with its potential to augment productivity and boost global economies. It is predicted that generative AI could add trillions of dollars in value to the global economy and significantly increase global GDP and productivity growth (McKinsey 2023, Goldman Sachs 2023). Generative AI's ability to generate a wide array of output forms—from text and code to images and videos—could potentially outpace human production capacity by 2030 (J.P. Morgan 2023). Particularly in marketing and sales, generative AI has the capacity to revolutionize the generation of creative content and reshape customer interactions, accounting for a substantial portion of the projected value of AI use cases (McKinsey 2023). Generative AI is already being used extensively in marketing for copywriting and creative generation. The partnership between technology titan Nvidia and the world's largest advertising agency, WPP, further exemplifies the transformative potential of generative AI in marketing (Ziady 2023). Given that marketing is dominant in shaping individuals' perceptions, we believe that Generative AI can have a significant impact on shaping society.

Furthermore, the impact of AI extends beyond commercial sectors, penetrating the realm of education. AI's capacity to personalize learning experiences across all educational levels has been recognized (Rudra 2023). Generative AI can improve the learning process by providing tailored information, stimulating creativity, and enhancing digital skills, preparing learners for future workplace demands. In the context of digital publishing in education, generative AI's potential for significant cost reduction of content creation (e.g., text, image, audio, video) underscores its practical benefits (Alphonso 2023).

While generative AI has numerous benefits in business and society, we must also remain aware and vigilant for potential drawbacks. Concerns arise about the technology's potential risks, including issues related to intellectual property rights, accuracy of output, explainability of results, and potential propagation of harmful biases. Of significant concern is the potential bias embedded in generative AI models. Different from traditional AI models that are often used for classification or prediction, generative AI models are used to create new content based on patterns from the training data, making it difficult to measure the bias as there is no single "correct" output. Instead, one would need to evaluate a range of generated content for patterns that reflect bias. Moreover, new content generated by these models such as visual content can directly shape users' perceptions, perpetuate harmful stereotypes, and even distort their beliefs, especially if the generated content is widely disseminated. For example, as educational tools increasingly harness generative AI, they inherently mold young minds and shape their worldview. Content tailored to student characteristics with potential bias may inadvertently propagate or solidify harmful stereotypes, shaping perceptions in ways that may be challenging to reverse.

Since generative AI models are often trained on vast quantities of data collected from the internet, lack of control over the sources presents a formidable challenge in auditing and updating the training data to handle potential bias. As the training data may encompass a plethora of perspectives, cultures, and ideologies, it becomes increasingly challenging to anticipate, let alone correct, the myriad biases that could inadvertently find their way into the model. Furthermore, the situation is complicated by the proprietary nature of several major generative AI models. These models, often owned and maintained by private entities, are not publicly available for scrutiny. This lack of transparency further exacerbates the problem, as the broader academic community and society cannot easily assess or rectify biases in such closed systems. Given their increasing popularity and rapid adoption in many domains, generative AI models may inadvertently reinforce detrimental biases and stereotypes in our society, particularly in marketing and education. The current lack of well-defined regulations and policies in the domain of generative AI presents further risks, creating space for potentially harmful applications.

This study investigates potential biases in three of the most popular text-to-image AI generators—Midjourney, Stable Diffusion, and DALL·E 2. These AI tools, ranging from commercial to open-source, can generate high-quality images from text prompts, a capability used by millions of users and deployed across various domains. Fig. 1 presents some examples of images generated by Midjourney with different textual descriptions.

Prompt: A Portrait of Personal Financial Advisors

Prompt: A Portrait of Chief Executives



Prompt: A Portrait of Dishwashers



Prompt: A Portrait of Anesthesiologist Assistants



**Fig. 1. Examples of Images Generated by Midjourney with Different Prompts.**

Using approximately 8,000 occupational portraits produced by these three AI generators, we demonstrate pervasive gender and racial biases across all three generative AI models. We find that women and Black individuals are significantly underrepresented in images generated by these tools. Moreover, when comparing this underrepresentation to different benchmarks such as BLS Labor Force Statistics (BLS 2022) and Google images, the disparity is even more pronounced than the *status quo*, intensifying the biases and stereotypes we are actively striving to rectify in today's society. Furthermore, troublingly, we find reinforcement of stereotypical gender portrayals in facial expressions and appearances. Women often depicted as more smiling and happier, while men are depicted as more neutral and angry, and perhaps more authoritative. While the magnitude of bias varies among the tools, the direction of bias was consistent across both commercial and open-source AI generators.

These biases may inadvertently perpetuate harmful societal stereotypes, with potentially far-reaching implications. For example, the overrepresentation of males in occupational images could deter aspiring female professionals, especially those at young ages when their minds are easily impressionable, thereby undermining efforts to promote Equality, Diversity, and Inclusion (EDI). Ideally, generative AI should serve all of humanity equally, shaping an inclusive future free from harmful biases. Our study underlines the critical need to identify and mitigate these biases, reinforcing the commitment to ensuring that generative AI inclusively benefits all of humanity.

## 2. Research Context and Data Generation

We utilized three popular text-to-image AI generators – Midjourney, Stable Diffusion, and DALL·E 2 – to generate images for various occupations. Each generator employs distinctive

models to transform text prompts into visual content. Notably, both Midjourney and DALL·E 2 are commercial AI image generators, while Stable Diffusion is an open-source program.

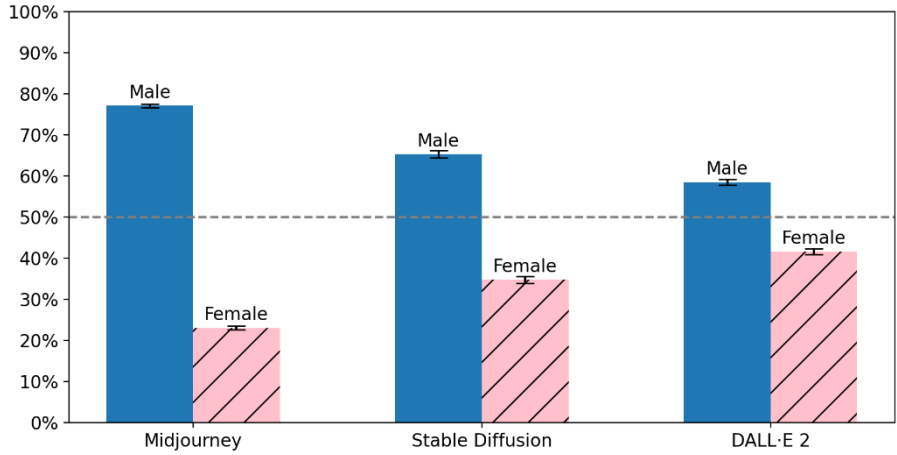
Midjourney was first published in July 2022 and is one of the leading drivers of emerging technology that uses AI to create visual imagery from text prompts (Salkowitz 2022), attracting approximately 15 million users (Broz 2023). Midjourney operates on a subscription-based model. DALL·E 2 is latest version of DALL·E, introduced by OpenAI in April 2022, and has amassed more than one million users (Bastian 2022). Its services are offered on a pay-per-use basis. Stable Diffusion, released for free in August 2022, utilizes a latent diffusion model to generate high-quality images from textual descriptions. As an open-source tool, Stable Diffusion has been used in pop music videos, Hollywood movies, and is used by more than 10 million people daily (Cai and Konrad 2023). The prowess of these generative AI models is well-recognized across diverse domains. For example, an AI-generated image was recently awarded the prize in an art competition, marking one of the first instances of AI-created artwork being so recognized (Roose 2022). This event signals a pivotal shift in societal and artistic perceptions and acceptance of AI-generated art.

To construct our data, we used consistent text prompts for all three models to generate images for occupations in the O\*NET database, which contains information on 1,016 occupations (O\*NET 2023). It is a regularly updated database of occupational characteristics and worker requirements information across the U.S. economy, representing one of the world’s largest, most comprehensive, widely used public repositories documenting detailed job-worker characteristics. Specifically, we structured our prompt as “A portrait of X” with X being each occupation in the O\*NET database. For DALL·E 2, which incurs a cost for each generated image, we produced two images per prompt. Similarly, with Midjourney, we followed this process to create corresponding occupational portraits. Since Midjourney automatically returns four images for every prompt, we obtained four images for each occupation. As for Stable Diffusion, an open-source tool, we used a pre-trained model called Realistic Vision, which is one of the best Stable Diffusion models for generating realistic images. We used this model to generate two images for every prompt. Consequently, about 8,000 different images were generated in total. We then conducted image analysis to investigate potential biases and stereotypes embedded in these AI-generated portrayals of various occupations.

### **3. Analyses and Results**


#### **3.1. Gender Bias**

In order to determine the gender distribution within our data, we utilized the Face++ API, one of the best face recognition APIs on the market (Eden AI 2023), to detect faces in each image and extract various features such as gender, smile, emotion, and age. Following this, we assessed the gender distribution within each image by calculating the percentage of women or men depicted, and subsequently computed an average across all images created by the same AI generators. As summarized in Fig. 3, the results revealed that all three AI image generators exhibited bias against women. Specifically, the representation of women in images created by Midjourney, Stable Diffusion, and DALL·E 2 was 23%, 35%, and 42%, respectively, all significantly lower than their male counterparts, which were represented at 77%, 65%, and 58%, respectively ( $p < 0.001$ ). While the degree of bias varied depending on the model, the direction of bias remained consistent across all three AI image generators, including both commercial and open-source models.



**Fig. 3. Gender Distribution in Occupational Portraits Created by AI Generators**

We classified each image into one of five exclusive categories: “Only Male” if it solely depicted males, “More Male” if the number of males exceeded females, “Equal” if there was an equal number of males and females, “More Female” if the females outnumbered males, and “Only Female” if it exclusively portrayed females. As summarized in Fig. 4, there was a significant skew towards male representation across all three AI generators. Notably, 52% of the images generated by Midjourney and 46% of the images generated by Stable Diffusion contained only males, indicating that approximately half of the generated occupational portraits by these models lacked any female representation.




**Fig. 4. Percentage of Different Images Created by AI Generators**

While the disparity in the generated images is alarming at an aggregate level, we examine how this disparity manifests across different types of jobs. To do this, we utilized a descriptor from the O\*NET database known as “Job Zone” as a proxy. A job zone categorizes occupations that require similar levels of education, related experience, and on-the-job training. Within the O\*NET database, there are 5 job zones, with Job Zone 1 requiring minimal preparation (3 months), and Job Zone 5 necessitating the most extensive preparation, typically over 4 years. Table 1 outlines the differences across job zones (Eloundou et al. 2023). It is observed that as the level of preparation required increases, so does the median income, with workers in Job Zone 1 earning \$30,230, and those in Job Zone 5 earning \$81,980.

Job Zone	Preparation Required	Education Required	Example Occupations	Median Income
1	None or little (0-3 months)	High school diploma or GED (optional)	Food preparation workers, dishwashers, floor sanders	\$30,230
2	Some (3-12 months)	High school diploma	Orderlies, customer service representatives, tellers	\$38,215
3	Medium (1-2 years)	Vocational school, on-the-job training, or associate’s degree	Electricians, barbers, medical assistants	\$54,815
4	Considerable (2-4 years)	Bachelor’s degree	Database administrators, graphic designers, cost estimators	\$77,345
5	Extensive (4+ years)	Master’s degree or higher	Pharmacists, lawyers, astronomers	\$81,980


**Table. 1. Occupations by Job Zone**

We then calculated the gender distribution in the portraits of occupations within each job zone. The findings, summarized in Fig. 5, indicate that gender bias permeates all job zones within the images generated by Midjourney and Stable Diffusion, with the extent of bias slightly decreasing as the required level of preparation increases. For DALL·E 2, the gender bias is evident across the first four job zones, with the representation of females surpassing that of males in Job Zone 5, which necessitates the highest level of preparation. A potential explanation for these findings is that jobs requiring lower levels of preparation might have exhibited stronger gendered representations historically. Such historical bias could be reflected in the data the models were trained on, leading the AI to generate images that not only mirror but potentially amplify these long-standing stereotypes. For example, recent data (BLS 2022) indicates an increasing presence of female workers across job zones: 39% in Job Zone 1, 38% in Job Zone 2, 43% in both Job Zone 3 and 4, and 59% in Job Zone 5.




**Fig. 5. Gender Distribution by Job Zone in Occupational Portraits Created by AI Generators**

Next, we investigate the variance in gender bias in relation to types of occupations using another proxy, an O\*NET occupation-level indicator called the “Bright Outlook.” Bright Outlook occupations are those expected to grow rapidly in the next several years, have large numbers of job openings, or are new and emerging occupations. Of the 1,016 occupations in the O\*NET database, 340 are classified as “Bright Outlook.” This indicator is positively correlated with the job zones (Pearson correlation coefficient = 0.131,  $p < 0.001$ ). We calculated the gender distribution in the portraits of Bright Outlook occupations, the results of which are depicted in Fig. 6. The gender bias was evident in all types of occupations across the three AI image generators, with a reduced bias in the portraits of Bright Outlook occupations.



**Fig. 6. Gender Distribution in Portraits of Bright Outlook Occupations**

To further investigate the variance in gender bias according to types of occupations, we used another O\*NET occupation-level indicator, “STEM,” denoting occupations that require education in science, technology, engineering, and mathematics (STEM) disciplines. Among the 1,016 occupations in the O\*NET database, 287 are labeled as “STEM.” This indicator is also positively correlated with the job zones (Pearson correlation coefficient = 0.563,  $p < 0.001$ ). We calculated the gender distribution in the portraits of STEM occupations, the results of which are summarized in Fig. 7. While the gender bias was still prevalent across all types of occupations for the three AI generators it was less pronounced in the portraits of STEM occupations.

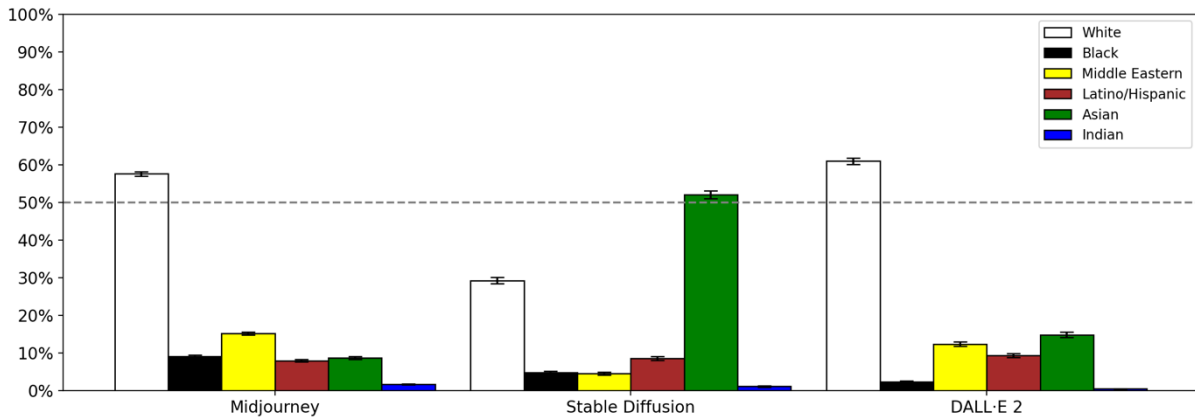


**Fig. 7. Average Gender Distribution in STEM Occupations**


### 3.2. Racial Bias

Next, we evaluated racial distribution by measuring the percentage of Black and White individuals portrayed in images produced by the three AI models. Utilizing the DeepFace framework (Serengil and Ozpinar 2020), we categorized the race of each face in the images as either White, Black, Asian, Latino, Indian, or Middle Eastern to compute the racial distribution in each image. We subsequently calculated the average value across all images produced by the same AI generators. Our findings, summarized in Fig. 8 and Fig. 9, reveal a bias against Black individuals across all three AI generators. Specifically, the percentage of Black individuals in the occupational portraits generated by Midjourney, Stable Diffusion, and DALL·E 2 was 9%, 5%, and 2%, respectively, all significantly lower than the percentages of White individuals ( $p < 0.001$ ). For Stable Diffusion, however, the representation of Asians exceeded 50%. This disproportion might be partially attributed to the training data used for the pre-trained Stable Diffusion model. Notably, for the two commercial AI generators, Midjourney and DALL·E 2, the representation of White individuals exceeded 50% consistently across all five job zones. However, the gaps between different races do not decrease as the required level of preparation increases, which was observed

in the gender analysis. This is likely because the historical bias against Black individuals does not diminish across job zones. For example, recent data (BLS 2022) shows a fluctuating presence of Black workers across job zones: 13% in Job Zone 1, 16% in Job Zone 2, 11% in Job Zone 3, and 10% in both Job Zone 4 and 5.



**Fig. 8. Average Racial Distribution in Images Created by AI Generators**



**Fig. 9. Average Racial Distribution by Job Zone in Images Created by AI Generators**

### 3.3. Benchmarking with BLS and Google Data

The gender and racial bias in the portraits of occupations generated by AI tools may stem from the inherent bias present in the training data, which reflects real-world disparities. To examine this, we obtained employment and demographic data from the BLS Labor Force Statistics (BLS 2022), from the Current Population Survey collected by the U.S. government, primarily capturing non-self-employed, documented workers in the formal economy. This data reported that the percentage of female workers of total employed is 46.8% (BLS 2022), significantly higher than the percentage of females in portraits of occupations created by Midjourney ( $p < 0.001$ ), Stable Diffusion ( $p < 0.001$ ), and DALL·E 2 ( $p < 0.05$ ); the percentage of Black workers is 12.6% (BLS 2022), significantly higher than that in occupational portraits created by Midjourney ( $p < 0.05$ ), Stable Diffusion ( $p < 0.001$ ), and DALL·E 2 ( $p < 0.001$ ). A comparison of their distribution is presented in Fig. 10. The results show that the gender and racial disparities in the real world are amplified in all three AI generative models. This is especially concerning given that the benchmark itself might encapsulate the societal bias that we are striving to counteract through various EDI initiatives.

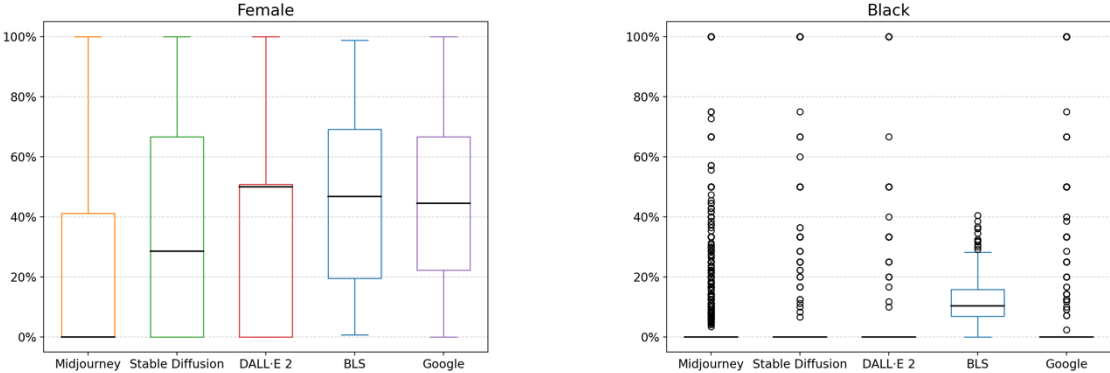
Furthermore, we examined occupation-specific disparities to pinpoint where the bias was most evident. Specifically, for each dataset, we calculated the disparities in female and Black representation relative to the BLS benchmarks. Occupations were then ranked based on these disparities. The results, presented in Fig. S1-S6 in the supplementary materials, reveal pervasive gender and racial bias across most occupations. Interestingly, results from the three models show some consistent patterns. For example, “food preparation and serving related workers” consistently appear among the top three occupations exhibiting the most pronounced gender bias against



women across all datasets, and “postal service clerks” rank at the top for the most pronounced racial bias against Black individuals in all three models. To assess how the generative AI models compare to publicly available data, we collected images from Google Image Search as another benchmark. To construct this benchmark, we used the same text prompts as search keywords in Google to collect images representing different occupations in the O\*NET database. For each search keyword, we saved the first ten images returned in the search results using Google’s API. This resulted in 10,160 images in total. These images are then analyzed to measure the gender and racial distribution in the data. Our results show that the percentage of females in the data collected from Google is 44.5%, which is not statistically different from that of BLS ( $p > 0.5$ ), but significantly higher than the percentage of females in portraits of occupations created by Midjourney ( $p < 0.001$ ), Stable Diffusion ( $p < 0.001$ ), and DALL·E 2 ( $p < 0.05$ ). This indicates that the image data on Google is representative of the gender distribution in the population and is perhaps not contributing to the gender bias in image generation models as far as gender is concerned.

The percentage of Black individuals in this Google data is 5.41%, which is significantly lower than that of BLS and Midjourney ( $p < 0.001$ ), not statistically different from that of Stable Diffusion, but significantly higher than the percentage of Black individuals in portraits of occupations created by DALL·E 2 ( $p < 0.001$ ). A comparison of the distribution is illustrated in Fig. 10, showing that all three AI generative models exhibited amplified gender and racial disparities in their generated images.

We argue that such gender and racial biases in representation are problematic, especially when these biases are amplified beyond real-world disparities and are worse than the status quo. For example, portraying primarily men may potentially dissuade the next generation of female and Black professionals and hinder efforts to promote EDI. Generative AI should benefit all of humanity and be shaped to be as inclusive as possible, at least not amplifying the biases in the status quo. Rather than reflecting, or even amplifying, the existing biases of today’s world, these tools should aspire to shape a better future that reflects equality and fairness.



**Fig. 10. Comparing Gender Distribution Across Datasets**

**3.4. Gender Stereotypes in Age, Smile, and Emotion**

It is important not only to consider who is shown in these images, but in which light they are shown. Facial expressions, age, approachability, and authority are various dimensions of these images that can impact how consumers perceive them. Drawing upon features such as gender, age, smile, and emotions extracted from detected faces, we gauged the average facial expressions and appearances for men and women in each image, and the outcomes are represented across the image generations. These results are illustrated in Fig. 11-12. Fig. 11 suggests that men appear older and

more senior than women in images generated by all three AI models which might inadvertently portray them in a position of authority as opposed to women that typically exhibit more smiles in those images. Moreover, women express more happiness (Fig. 12) than men across all three AI generators. Men, on the other hand, typically express more neutral emotion and anger than women across all three AI generators. This stereotype persists across all three generative models.




Fig. 11. Average Age and Smile of Males vs Females in Images Generated by AI




Fig. 12. Average Emotions by Gender in Images Created by AI Generators

4. Conclusions

This study highlights the presence of various biases and stereotypes embedded within AI image generators. Our research underscores the urgency of addressing these issues, as these AI tools are increasingly employed in diverse sectors including marketing and education. The potential fallout of such biases, including the shaping and reinforcement of societal perceptions and stereotypes, warrants our meticulous attention and exploration.

By illuminating how AI interprets different occupations, our research contributes to a crucial understanding of the bias endemic in AI systems. The degree of bias pervasive within these generative AI tools is a matter of significant concern, calling for immediate introspection and remedial action within the AI community.

As our reliance on AI intensifies, it becomes imperative to ensure fairness and equity in the design and deployment of these technologies. We must prioritize the development of generative AI systems that are not only technologically advanced but are also shaped by an ethical commitment to inclusivity and equity. This ensures these powerful tools are designed to benefit all users equitably, fostering a more inclusive future.

References

Alphonso, G. (2023). Generative AI: Education in the age of innovation. *Forbes*. <https://www.forbes.com/sites/forbestechcouncil/2023/03/03/generative-ai-education-in-the-age-of-innovation/>

- Bastian, M. (2022). DALL-E 2 has more than one million users, new feature released. *The Decoder*. <https://the-decoder.com/dall-e-2-has-one-million-users-new-feature-rolls-out/>
- BLS. (2022). Labor Force Statistics from the Current Population Survey. <https://www.bls.gov/cps/cpsaat11.htm>
- Broz, M. (2023). Midjourney statistics (Updated July 2023). *Photutorial*. <https://photutorial.com/midjourney-statistics/>
- Cai, K., & Konrad, A. (2023). Six things you didn't know about ChatGPT, Stable Diffusion and the future of generative AI. *Forbes*. <https://www.forbes.com/sites/kenrickcai/2023/02/02/things-you-didnt-know-chatgpt-stable-diffusion-generative-ai/?sh=149feaa2b5e3>
- Eden AI. (2023). Best face recognition APIs in 2023. <https://www.edenai.co/post/best-face-recognition-apis>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv preprint arXiv:2303.10130.
- J.P. Morgan. (2023). The rise of generative AI. <https://www.jpmorgan.com/insights/research/generative-ai>
- Liao, S. (2023). A.I. may help design your favorite video game character. *The New York Times*. <https://www.nytimes.com/2023/05/22/arts/blizzard-diffusion-ai-video-games.html>
- McKinsey. (2023). The economic potential of generative AI: The next productivity frontier. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/The-economic-potential-of-generative-AI-The-next-productivity-frontier#introduction>
- O\*NET. (2023). <https://www.onetonline.org/find/all>
- Roose, K. (2022). AI-generated art won a prize. Artists aren't happy. *The New York Times*. <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>
- Rudra, S. (2023). Could the disruptive power of AI-generated images and videos impact K–12 learning? Technology solutions that drive education. *EdTech*. <https://edtechmagazine.com/k12/article/2023/03/could-disruptive-power-ai-generated-images-and-videos-impact-k-12-learning>
- Salkowitz, R. (2022). Midjourney founder David Holz on the impact of AI on art, imagination and the creative economy. *Forbes*. <https://www.forbes.com/sites/robsalkowitz/2022/09/16/midjourney-founder-david-holz-on-the-impact-of-ai-on-art-imagination-and-the-creative-economy/?sh=1b0cd8c02d2b>
- Serengil, S. I., & Ozpinar, A. (2020). Lightface: A hybrid deep face recognition framework. In *2020 Innovations in Intelligent Systems and Applications Conference (ASYU)* (pp. 1-5). IEEE.
- Ziady, H. (2023). The world's biggest ad agency is going all in on AI with Nvidia's help. *CNN*. <https://www.cnn.com/2023/05/29/tech/nvidia-wpp-ai-advertising>



## Supplementary Materials

Fig. S1. Disparity in Female Ratios: Midjourney vs BLS Data by Occupation




Fig. S2. Disparity in Female Ratios: Stable Diffusion vs BLS Data by Occupation




Fig. S3. Disparity in Female Ratios: DALL·E 2 vs BLS Data by Occupation




Fig. S4. Disparity in Black Ratios: Midjourney vs BLS Data by Occupation






Fig. S5. Disparity in Black Ratios: Stable Diffusion vs BLS Data by Occupation




Fig. S6. Disparity in Black Ratios: DALL·E 2 vs BLS Data by Occupation

