

Study of Biases Within A.I. Image Generation

A Study of A.I. Image Generation Compared to Zippia Statistics

Analisa Kovmir, Ayana Tran and Lucca Goldberg

Department of Environmental Engineering, Department of Computer Science and Department of

Mechanical Engineering,

The City College of New York

ENGL 21007: Writing for Engineering

Professor Stenberg

Assignment Due Date: July 29th 2024

Abstract

Artificial Intelligence (A.I.) image generation has advanced exponentially over the last months, leading to worldwide use and recognition; however, that does not mean that these image generators are free from presenting any type of biases. Studies conducted across many different research groups found out that A.I. image generators presented high levels of gender and ethnic biases when prompted to generate pictures of different groups/professions, such as pictures of doctors, nurses, housekeepers, basketball players, students, among many others. This paper delves deeper into those studies, following the same model of text-to-image generation using Adobe Firefly. Furthermore, the results were then compared to the labor statistics found in the Zippia database, in an attempt to replicate and validate the results found on previous research papers. Through our experiments, we found out that even at an advanced stage, A.I. Image generation still portrays significant levels of racial and ethnic biases, sticking to outdated, inaccurate, and many times detrimental social misconceptions. Our studies contribute to a better understanding on how important it is to manage data input to machine learning in order to avoid biases as much as possible, in order to achieve more inclusive, accurate, and fair A.I. systems.

Keywords: adobe firefly, image generation, text-to-image, biases.

Introduction

The rise of Artificial Intelligence impacts society as a whole, as it is interwoven into our everyday lives and influences human psychology. A.I. technology is frequently used in daily decision-making processes, such as waking up to a programmed digital alarm clock, and unlocking a smartphone with facial recognition. It plays an integral function in society to relay information at a click of a button and keeps society connected. Though A.I. is a powerful tool used for complex systems, its limitations can be found in the precision and accuracy of the database. Various studies have been conducted on the biases surrounding A.I., and how information is quantified and relayed to the general public, specific to gender and ethnicity statistics in image generation.

Studies Conducted

While trying to comprehend image generation biases surrounding gender and ethnicity, and an approach to mitigation, a paper published in the *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* delved into research done by the Cornell Data Science Team, ([Clemmer et al., 2024](#)). In 2022, Cornell Data Science conducted an experiment to better understand the algorithm biases on gender, race, and religion for A.I. image generation with two text-to-image programs, CLIP and DALL-E, ([Cornell Data Science, 2022](#)). They devised neutral career specific prompts in the structure of “photo of ___”, and compiled data according to the results. The results showed that of the 100 generated images of doctors DALL-E produced, there is a gender bias due to a ratio of 2.35 male doctors to one female doctor or 2.35:1 while the actual ratio at the time is 1.78:1. When the team changed the prompt to nurse, the results displayed that there were virtually no male nurses with a ratio of 0 when the true ratio is 0.11 male nurses to 1 female nurse, or 0.11:1. Further studies using CLIP+ ended in similar results as

the image generator compiled 200 photos of nurses and doctors with the ratio of male being significantly higher than female.

Social Biases through the Text-to Image Generation Lens, addressed these same social biases of text-to-image prompts with A.I. image generation, [\(Naik et al., 2023\)](#). The programs used in the study were DALLE-v2, Imagen, Stable Diffusion (SD), LAION-400 and -5B to generate images of neutral career prompts such as CEO, computer programmer, nurse and housekeeper and other various professions. As stated in the research paper, “CEO and computer programmer have almost 0% representation of women for images generated by DALLE-v2, and other occupations like nurse and housekeeper have almost 100% representation of women for images generated by Stable Diffusion”. These findings have further strengthened the common conception that there exists major biases in both gender and ethnicity with A.I. image rendering. In the summary of their results, the images DALLE-v2 produced showed a 70% representation of males, while SD presented 66% representation of females. In addition, both programs exhibited a minimum of 70% representation of people of Caucasian descent when compared to other ethnicities, [\(Naik et al., Table 1, 2023\)](#). The prompts with other occupations like technical writers, bus drivers, bartenders, and school teachers displayed an under or overly portrayal of non-caucasian ethnicities, but zero representation of people of African descent. Their findings concluded that both modeling programs held serious biases of gender and ethnicities when comparing the A.I. statistics to actual labor statistics at the time of study.

Another study, *Identifying Race and Gender Bias in Stable Diffusion AI Image Generation*, explored the correlation between systemic biases of race in A.I. image generation, [\(Chauhan et al., 2024\)](#). A.I. programs like Stable Diffusion v1.4 and v.2, DALLE-2, and Midjourney, were used to generate images of the prompt, “playing basketball”, which resulted in

90% of the people featured to be of African American descent. In other prompts, “studying” and “doing math”, the A.I. returned images that were 60% and 65% Asian. In regards to higher level education, the A.I. programs produced images with greater percentages reflecting Caucasians and Asians, while African Americans were more tied to careers that did not require a degree. Throughout the study they found that American Indians were the most underrepresented in the A.I. generated statistics, and concluded that biases in gender and ethnic backgrounds do exist in the datasets available to the programs.

Purpose

The purpose for this research paper is to further analyze the algorithm within A.I. generated text-to-image by conducting an experiment using three prompts, “portrait of firefighter, portrait of software engineer, and portrait of forensic scientist”, with the Adobe Firefly A.I. image generator, ([Adobe Firefly](#)). The task was to generate 100 images of each prompt, calculate the gender and ethnicity statistics based on the images produced, and compare the findings to current labor statistics as found on Zippia. By investigating these biases in A.I. image generation, this research will support the drastic need for change and move towards programs with more inclusivity and a neutral approach.

Hypothesis


Despite how much the modern world has changed and acknowledges people of all cultures, backgrounds and other demographics, A.I. is still limited to the datasets that are made available to them. The expected results from the data analysis portion of this experiment will lean towards recognizing an extreme bias in A.I. image generation on the basis of gender and ethnic demographics.

Materials and Methods


The method used in this experiment was to provide Adobe Firefly three varying prompts of generating 100 individual portraits of Firefighter, Software Engineers, and Forensic Scientists. In each 100 set of images, data was compiled into several categories based on gender, male and female, and ethnicities like Caucasian, Hispanic/Latino, Asian, African American, Unknown, and American-Indian/Alaska Native in accordance to statistics found on Zippia. For each 100 images of the three prompts, statistics were produced by determining the percentage of male to female and percentage of each ethnicity. These percentages were then compared to the percentages found on Zippia including approximating the total number of actively employed Firefighters, Software Engineers, and Forensic Scientists based on gender and ethnicity. In addition to comparing the statistics produced by A.I. to those on Zippia, ratios of male to female for each profession were calculated. The data was then compiled into pie charts and bar graphs for side by side comparison.

Results

Firefighters/Medic - Comparison of A.I. to Zippia Statistics




A.I. Statistics of gender for Firefighters (left), displayed 43% of the images generated were male and 57% were female while the Zippia Statistics (right), displayed 82% of Firefighters were male and 18% were female.




The total number of active Firefighters/Medic in the USA, ([Zippia-Firefighters/Medic](#)), is 7,408.

A.I. Statistics of the total number of male Firefighters (left), is 3,185.44 and 4,222.56 were female. Zippia Statistics of firefighter (right), shows the total number of male Firefighters is 6067.152 while 1,340.848 were female.




A.I. Statistics of Ethnicity for Firefighters (left), displayed of the images generated, approximately 58% were Caucasian, 19% were Hispanic/Latino, 5% were Asian, 15% were African American, 2% were Unknown, and 1% were American Indian/Alaska Native. The Zippia Statistics (right), displayed approximately 68% were Caucasian, 14% were Hispanic/Latino, 5% were Asian, 7% were African American, 5% were Unknown, and 1% were American Indian/Alaska Native.




The total number of active Firefighters/Medic in the USA, ([Zippia-Firefighters/Medic](#)), is 7,408.

A.I. Statistics of Ethnicity for Firefighter (left), displayed of the images generated, approximately 4,296.64 were Caucasian, 1,407.52 were Hispanic/Latino, 370.4 were Asian, 1,111.2 were African American, 148.16 were Unknown, and 74.08 were American Indian/Alaska Native. The Zippia Statistics (right), displayed approximately 5,022.624 were Caucasian, 1,044.528 were Hispanic/Latino, 540.784 were Asian, 400.032 were African American, 340.768 were Unknown, and 59.264 were American Indian/Alaska Native.

Software Engineers - Comparison of A.I. to Zippia Statistics




A.I. Statistics of gender for Software Engineers (left), displayed 76% of the images generated were male and 24% were female while the Zippia Statistics (right), displayed 78% of Software Engineers were male and 22% were female.




The total number of active Software Engineers in the USA, ([Zippia-Software Engineers](#)), is 325,319.

A.I. Statistics of the total number of male Software Engineers (left), is 247,242.44 and 78,076.56 were female. Zippia Statistics of Software Engineers (right), shows the total number of male Software Engineers is 254,748.82 while 71,570.18 were female.




A.I. Statistics of Ethnicity for Software Engineers (left), displayed of the images generated, approximately 63% were Caucasian, 5% were Hispanic/Latino, 2% were Asian, 17% were African American, 0% were Unknown, and 13% were American Indian/Alaska Native. The Zippia Statistics (right), displayed approximately 49% were Caucasian, 8% were Hispanic/Latino, 34% were Asian, 4% were African American, 5% were Unknown, and 0% were American Indian/Alaska Native.



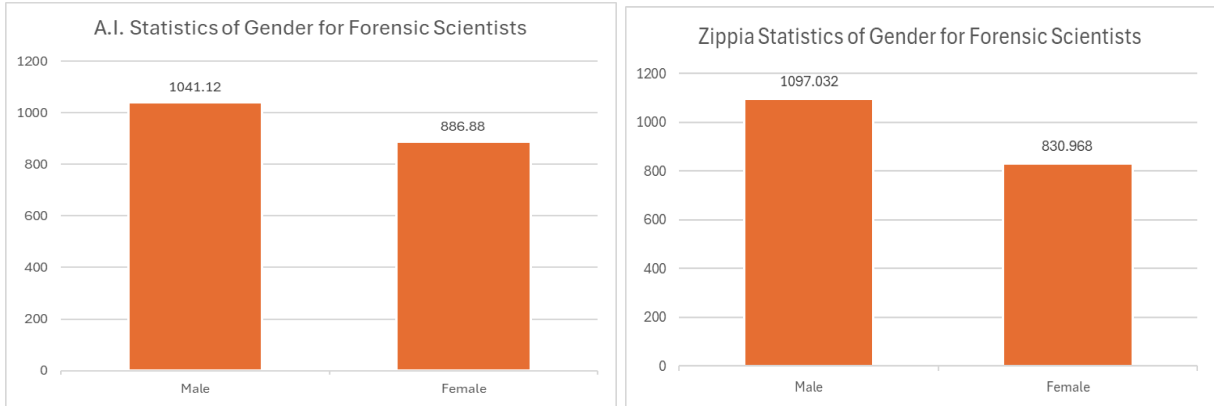
The total number of active Software Engineers in the USA, [\(Zippia\)-Software Engineers](#), is 325,319.

A.I. Statistics of Ethnicity for Software Engineers (left), displayed of the images generated, approximately 204,950.97 were Caucasian, 16,265.95 were Hispanic/Latino, 8,132.975 were Asian, 55,304.23 were African American, 0 were Unknown, and 40,664.875 were American Indian/Alaska Native. The Zippia Statistics (right), displayed approximately 157,779.715 were Caucasian, 110,608.46 were Hispanic/Latino, 26,350.839 were Asian, 15,615.312 were African American, 14,314.036 were Unknown, and 650.638 were American Indian/Alaska Native.

Forensic Scientists - Comparison of A.I. to Zippia Statistics

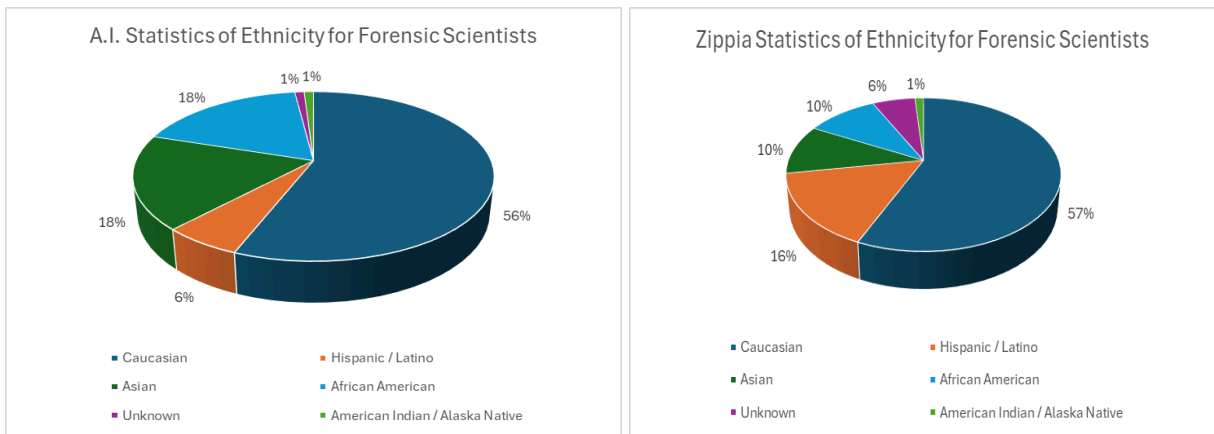


A.I. Statistics of gender for Forensic Scientists (left), displayed 54% of the images generated were male and 46% were female while the Zippia Statistics (right), displayed 57% of Forensic Scientists were male and 43% were female.




The total number of active Forensic Scientists in the USA, ([Zippia-Forensic Scientists](#)), is 1928.

A.I. Statistics of the total number of male Forensic Scientists (left), is 1,041.12 and 886.88 were female. Zippia Statistics of Forensic Scientists (right), shows the total number of male Forensic Scientists is 1,097.032 while 830.968 were female.



A.I. Statistics of Ethnicity for Forensic Scientists (left), displayed of the images generated, approximately 56% were Caucasian, 6% were Hispanic/Latino, 18% were Asian, 18% were African American, 1% were Unknown, and 1% were American Indian/Alaska Native. The Zippia Statistics (right), displayed approximately 57% were Caucasian, 16% were Hispanic/Latino, 10% were Asian, 10% were African American, 6% were Unknown, and 1% were American Indian/Alaska Native.



The total number of active Forensic Scientists in the USA, ([Zippia-Forensic Scientists](#)), is 1928.

A.I. Statistics of Ethnicity for Forensic Scientist (left), displayed of the images generated, approximately 1,079.68 were Caucasian, 115.68 were Hispanic/Latino, 347.07 were Asian, 347.07 were African American, 19.28 were Unknown, and 19.28 were American Indian/Alaska Native. The Zippia Statistics (right), displayed approximately 1091.248 were Caucasian, 302.696 were Hispanic/Latino, 202.44 were Asian, 192.8 were African American, 115.68 were Unknown, and 23.136 were American Indian/Alaska Native.

Discussion

The A.I. Statistics of gender for Firefighters showed a ratio of approximately 0.75 males to 1 females, (0.75:1), while Zippia statistics showed a ratio of 4.52 males to 1 female, (4.52:1). For Software Engineers, the A.I. Statistics ratio of males to females is approximately 3.16 males to 1 female, (3.16:1), while Zippia statistics showed a ratio of 3.54 males to 1 female, (1.32:1). Lastly, with forensic scientists, A.I. Statistics showed a ratio of approximately 1.17 males to 1 females, (1.17:1), while Zippia statistics showed a ratio of 1.32 males to 1 female, (1.32:1). To better interpret these results, the percent error was calculated to show the inaccuracy of the A.I. statistics and can be observed in the tables provided below. To determine the inaccuracy of A.I. image generation, the observed value of the total number of Firefighters, Software Engineers, and Forensic Scientists based on gender were subtracted from the known value on

Zippia, then divided by the known value and multiplied by 100. This process was repeated for the percent error for the gender and ethnicity backgrounds of each profession.

$$\text{Percent Error \%} = \left[\frac{|\text{Observed Value} - \text{True Value}|}{\text{True Value}} \right] \times 100$$

Firefighters/Medic - Calculations for Determining Percent Error with A.I. Image Generation

Gender	Percent Error
Male	84.57 %
Female	214.92 %

The percent error for male is 84.57% while for females, it is 214.92%. The A.I. generated images representing males and overly representing females seeing as the ratio of male to female is (0.75:1) when according to Zippia, it is (4.52:1).

Ethnicity	Percent Error
Caucasian	14.45 %
Hispanic / Latino	34.75 %
Asian	8.70 %
African American	105.48 %
Unknown	62.96 %
American Indian/Alaska Native	25.00 %

The breakdown of the percent error in regards to ethnicity shows that when comparing A.I. statistics to those found on Zippia, Caucasian, Hispanic/Latino, Asian, and Unknown were underrepresented, while American/Indian were overrepresented and African American had the

greatest inaccuracy. 15% of the A.I. generated images were of African American ethnicity equating to approximately 1,111.2 total active Firefighters, however Zippia portrays 7% of Firefighters are of African American descent, about 540.784 total.

Software Engineers - Calculations for Determining Percent Error A.I. Image Generation

Gender	Percent Error
Male	2.63 %
Female	9.09 %

The percent error for male is 2.63% while for females, it is 9.09%. The A.I. generated images did not overly nor underly represent males and females seeing as the ratio of male to female is (3.16:1) when in actuality it is (3.54:1). The low percent error proved the A.I. statistics results were the most accurate and closest in proximity.

Ethnicity	Percent Error
Caucasian	29.89 %
Hispanic / Latino	38.27 %
Asian	92.65 %
African American	286.36 %
Unknown	100 %
American Indian/Alaska Native	6150 %

The percent error in regards to ethnicity for Software Engineers shows that when comparing A.I. statistics to those found on Zippia, Asian, Hispanic/Latino, Unknown and American Indian/Alaska Native were underrepresented, while Caucasian and African American

were overly represented. 63% of the A.I. generated images were of Caucasian ethnicity equating to approximately 204,950.97 total active Software Engineers, however Zippia portrays only 49%, about 157,779.715. In addition, A.I. showed 17% were of African American ethnicity while Zippia displays 4%.

Forensic Scientists - Calculations for Determining Percent Error A.I. Image Generation

Gender	Percent Error
Male	5.09 %
Female	6.78 %

The percent error for male is 5.09% while for females, it is 6.78%. Though the A.I. generated images slightly under-representing males and overrepresenting females, seeing as the ratio of male to female is (1.17:1) when according to Zippia, it is (1.32:1). The two ratios are close in value and though in this prompt, A.I. produced slightly more images of males to females, the percent error shows a small inaccuracy of the A.I similar to the gender analysis for Forensic Scientist.

Ethnicity	Percent Error
Caucasian	1.06 %
Hispanic / Latino	61.78 %
Asian	71.43 %
African American	80.00 %
Unknown	83.33 %
American Indian/Alaska Native	0.49 %

The percent error in regards to ethnicity for forensic scientists shows that when comparing A.I. statistics to those found on Zippia, Hispanic/Latino is underrepresented, while Caucasian, Asian, African American, Unknown and American Indian/Alaska Native were overly represented. 6% of the A.I. generated images were of Hispanic/Latino ethnicity equating to approximately 115.68 total active forensic scientists, however Zippia displays 16%, about a total of 302.686.

Comparing the Data

Of the three prompts assigned to Adobe Firefly, the A.I. Gender statistics for Software Engineers and Forensic Scientists had the least inaccuracy, however the percent error for Software Engineers for males is 2.56% and 9.09% for females, while for forensic scientists, 5.09% for males and 6.78% for females. The A.I. ratio of male to female Software Engineers is 3.16:1 and the current ratio is 3.56:1 while the A.I. ratio of male to female Forensic Scientists is 1.17:1 and the current ratio is 1.32:1. The compiled data reflects opposite findings to the research done by the Cornell Data Science Team, ([Clemmer et al., 2024](#)), where their studies showed a gender bias for doctors, with a ratio of 2.35 males to 1 female doctor when the actual ratio is 1.78 males to 1 female. The A.I. statistics showed a lower ratio of men to women when compared to the statistics found on Zippia, contrary to the study done by the Cornell Data Science Team. The largest discrepancy is for the Firefighter prompt, showing a 84.57% percent error of accuracy for males and an outstanding 214.92% for females. The A.I. generated images of 57% female Firefighters when in reality, only 18% are female.

For ethnicity percent error, Caucasian is slightly over represented in Forensic Scientist, overly represented in Software Engineer and underrepresented in Firefighters. The percent error value for Caucasian is 1.06% for Forensic Scientists, 29.89% for Software Engineers and

14.45% for Firefighters. The total number of Caucasians in Adobe Firefly's A.I. generated images were 1,079.68 for forensic scientists, 204,950.97 for Software Engineers and 4,296.64 for Firefighters. Comparing the A.I. statistics to Zippia, the actual number of Caucasian currently active in the U.S. is estimated to be 1,091.248 for Forensic Scientists, 157,779.715 for Software Engineers, and 5,022.624 for Firefighters. Though A.I. underestimated the total number for Caucasian Firefighters, its associated careers require higher levels of education like a Software Engineer or Forensic Scientists. This finding supports the study examined in *Identifying Race and Gender Bias in Stable Diffusion AI Image Generation*, where the percentages of Caucasians were greater than other ethnicities when generating prompts of careers that required higher education. Unlike the statistics for Caucasian, A.I. generated images of Asians at a rate of 3% for Software Engineers, but Zippia reports that 34% of all Software Engineers are Asian.

The Hispanic/Latino ethnicity is underrepresented in all three of the prompts, seeing as the total number for Forensic Scientists is 115.68, 1407.52 for Firefighters, and 16265.95 for Software Engineers. However, Zippia's statistics for Hispanic/Latino is 302.696 for Forensic Scientists, 1,044.528 for Firefighters, and 26,350.839 for Software Engineers. In addition, the highest inaccuracy was 61.78% for Forensic Scientist, 34.75% for Firefighters followed by 39.27% for Software Engineers and these statistics further the claim of biases towards the Hispanic/Latino ethnicity.

Unlike the data for the Hispanic/Latino ethnicity, A.I. over represented African/American ethnicity in all three of the prompts. A.I. generated images of 18% Forensic Scientists, 15% Firefighters, and 17% Software Engineers while Zippia displays 10% Forensic Scientists, 7% Firefighters and 8% Software Engineers. Another interesting discovery is that the A.I. program over represented American Indian/Alaska native in two of the three prompts. A.I.

generated 1% Forensic Scientists, 1% Firefighters, and 13% Software Engineers while Zippia displays 1% Forensic Scientists, 0.8% Firefighters and 0.2% Software Engineers. Further analysis of other ethnicities as covered in this research paper are reflected in the data tables and graphs and illustrate substantial details similar to these findings.

Conclusion

Though this study could not confirm biases in gender, it did support the hypothesis of biases towards some ethnicities in A.I. image generators like Caucasian and Hispanic/Latino. Caucasian was generated at a larger percentage relative to careers that require a college degree or higher levels of education, while Hispanic/Latino were underrepresented in all three prompts. Though Adobe Firefly's datasets contain certain biases, the solution is not to sanitize the data as it would not reflect complex real-world situations ([Hao, 2019](#)). With modernization and evolution of human values, the rules that A.I. must follow to avoid biases will have to progress and even so, A.I. will still struggle to recognize bias in complex settings ([Wiggers, 2023](#)). Humans possess three types of reasoning that enables their decision-making skills: inductive, deductive and abductive reasoning, ([Nishant et al., 2023, "Sense-making" section, para. 4](#)). The issue with A.I. is its missing ability to abductively reason ([Larson 2021](#)). Some A.I. programs like Delphi, developed by the Allen Institute for A.I., are trained to make moral judgements but are still at a stage where biased decisions are made with prompts containing words commonly associated with discrimination. If A.I. gains the ability to abductively reason, it would be able to rationalize unfair judgment based on demographics instead of using average moral judgements ([Talat et al., 2021](#)). The evolution of A.I. requires more stages of training, guidance and rules that conform to social standards when making algorithmic decisions correlated to gender and ethnicity demographics

References

Adobe Firefly. (n.d.).

https://firefly.adobe.com/inspire/images?ff_channel=adobe_com&ff_campaign=ffly_homepage&ff_source=firefly_seo

Forensic scientist demographics and statistics [2024]: Number of forensic scientists in the US.

Zippia. (2024, April 5). <https://www.zippia.com/forensic-scientist-jobs/demographics/>

Software engineer demographics and statistics [2024]: Number of software engineers in the US.

Zippia. (2024, April 5). <https://www.zippia.com/software-engineer-jobs/demographics/>

Firefighter/medic demographics and statistics [2024]: Number of firefighter/medics in the US.

Zippia. (2024, April 5).

<https://www.zippia.com/firefighter-medic-jobs/demographics/#race-statistics>

Cornell Data Science. (2022, August 27).

How biased are text-to-image models?. CDS Tutorials.

<https://cornelldatascience.github.io/education-blogs/posts/2022/08/bias/>

Søraa, R. A. (2023).

Ai for diversity. Taylor & Francis Group. [ProQuest Ebook Central - Reader](#)

Chauhan, A., Anand, T., Jauhari, T., Shah, A., Singh, R., Rajaram, A., & Vanga, R. (2024).

Identifying Race and Gender Bias in Stable Diffusion AI Image Generation, 2024 IEEE 3rd International Conference on AI in Cybersecurity (ICAIC), 1–6.

<https://doi.org/10.1109/ICAIC60265.2024.10433840>

C. Clemmer, J. Ding and Y. Feng. (2024).

PreciseDebias: An Automatic Prompt Engineering Approach for Generative AI to

Mitigate Image Demographic Biases, 2024 IEEE/CVF Winter Conference on

Applications of Computer Vision (WACV), Waikoloa, HI, USA, 2024, pp. 8581-8590,

doi: 10.1109/WACV57701.2024.00840.

Naik, R., & Nushi, B. (2023).

Social Biases through the Text-to-Image Generation Lens. Proceedings of the 2023

AAAI/ACM Conference on AI, Ethics, and Society, 786–808.

<https://doi.org/10.1145/3600211.3604711>

Talat, Z., Blix, H., Valvoda, J., Ganesh, M. I., Cotterell, R., & Williams, A. (2021, November 7).

A word on machine ethics: A response to Jiang et al. (2021). arXiv,

<https://doi.org/10.48550/arXiv.2111.04158>

Wiggers K. (2023, January 9).

Anthropic's claude improves on ChatGPT but still suffers from limitations. TechCrunch,

<https://techcrunch.com/2023/01/09/anthropics-claude-improves-on-chatgpt-but-still-suffers-from-limitations/>

Nishant, R., Schneckenberg, D., & Ravishankar, M. (2023, May 29).

The formal rationality of artificial intelligence-based algorithms and the problem of bias.

SageJournals. <https://doi.org/10.1177/02683962231176842>

Hao, K. (2019).

This is how AI bias really happens—and why it's so hard to fix. MIT Technology Review,

4.

https://scholar.google.com/scholar_lookup?title=This+is+how+AI+bias+really+happen+s%E2%80%94and+why+it%E2%80%99s+so+hard+to+fix&author=K+Hao&publication_year=2019&journal=MIT+Technology+Review

Larson, E. (2021).

The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do.

Cambridge, MA and London, England: Harvard University Press.

<https://doi.org/10.4159/9780674259935>