

Evaluation of Gender, Age, and Race Bias within AI

Ryan Ruiz, Alexander Michuad, Raghvan Mahajan

ENGL 21007

Prof. Pamela Stenberg

November 12, 2023

Table of Contents

Page 2.....	Table of Contents
Page 3.....	Abstract
Page 3.....	Introduction
Page 3.....	Methodology
Page 5.....	Results
Page 10.....	Discussion
Page 13.....	Conclusion
Page 14.....	References

Abstract

In this lab report, we studied the websites that generate images with A.I. software that portray people in certain career fields with any stereotypes or biases. Our group chose three gender-neutral occupations and observed the race, gender, and age of the images that the AI generated. The three occupations that we picked were writer, engineer, and firefighter. Two of the group members used Microsoft Bing Image Creator to find images with the search terms firefighter and engineer. Our third group member chose Craiyon v3 to generate images related to writers. After getting 100 photos for each profession, we examined each photo to see the person's race, gender, and age. Our group compiled all the photos together and put the gender, age, and race data in pie charts. After looking up the real-life statistics on Zippa, we compared them with our results. We then discussed how accurate the AI image generator is in depicting the gender, race, and age with the job demographics and we see how bias is shown.

Introduction

Stereotypes and biases have been made on what gender, age, and race group may have a certain job or career. As people's understanding of the world evolved, so did our technology. As the human race as a whole enters the 21st century, our technology continues to evolve. We have been able to create devices that give us access to a wealth of never-ending knowledge. In more recent years, technology has continued to evolve and has gained intelligence. This comes in the form of Artificial Intelligence (AI). AI is capable of almost anything. From creating images based on text to being able to create a recipe for a five-star meal.

As AI continues to be developed, it becomes more human in how it interacts with us and answers our questions. In previous papers revolving around this topic, many have looked into

algorithms that are used to create AI results containing gender bias. Other papers show how AI has bias within image recognition.

However, these previous papers mostly discuss how this problem is prevalent within our society or incorporate jargon within their paper. This tends to limit the audience of the paper as not everyone is well-versed in AI algorithms. This is apparent in “Detecting Race and gender bias in visual representation of AI on web search engines”, and “AI Ageism: A Critical Roadmap for studying age Discrimination and Exclusion in digitized societies”. They show relevant data but overwhelm the reader with an abundance of information about the topic or information that most people won't understand.

However, this paper discusses how this gender, age, and race bias is relevant in text-to-image AI generation. This method compared to other papers is much easier for readers to understand and does not include as much jargon. Since people will always naturally have stereotypes about what gender, age, and race are appointed to a certain role, that same bias tends to leak into the algorithms that are used to create the images. The purpose of the study is to determine whether this is true. If there is gender, age, and race bias within AI, then the majority of the AI-generated images received should display unrealistic standards for each career choice.

Methodology

To collect data, we utilized a set of A.I. image generator tools available to the public on the Internet. They were Microsoft Bing Image Creator and Craiyon v3. We then collected a pool of 300 images, 200 from Bing and 100 from Craiyon v3. We used simple, asexual prompts in the A.I. generators, only involving the name of the profession chosen: writer, firefighter, and

engineer. We used Bing to create firefighter and engineer images and Craiyon for the writer images. After generating the A.I. images, they were each classified by gender, age, and race. The images were classed in different categories. The gender classifications were: female, male, and neutral. The age classifications were: 18-30 and 30-60. The race classifications differed between Black, White, Asian, and Indian. We then used our visual judgment to decide whether the people depicted in the images were male or female and their ages. The images where gender was not identifiable were labeled as neutral. The images where the face was either hidden or not visible were also labeled as neutral.

Results

After generating all 300 images and classing them by gender, age, and race, we gathered the following main idea. Most people depicted in the generated images were depicted as young. Bing especially generated images with almost entirely people between 18 and 30 years old. The biggest example of this is the percentage of engineering women. Out of 100 engineering females, 100% were between 18 and 30. As for the males, 28.6% were between 18 and 30 years old while 71.4% were between 30 and 60.

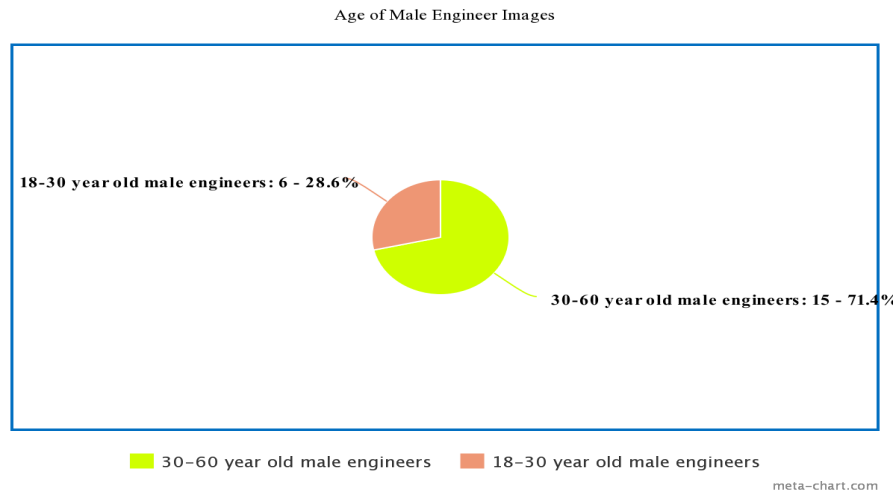


Figure 1. Age of Male Engineer images

For the gender results, Bing mostly generated images representing women. Almost 80% of the people depicted as engineers were women, with only 22% being men and only 1% with people without noticeable features.




Figure 2. Gender of Engineer Images

For the race results, Bing generated mostly White and Asian engineers with 42% and 32% respectively. The other two races included were Black with 21% and Middle Eastern with 5%.

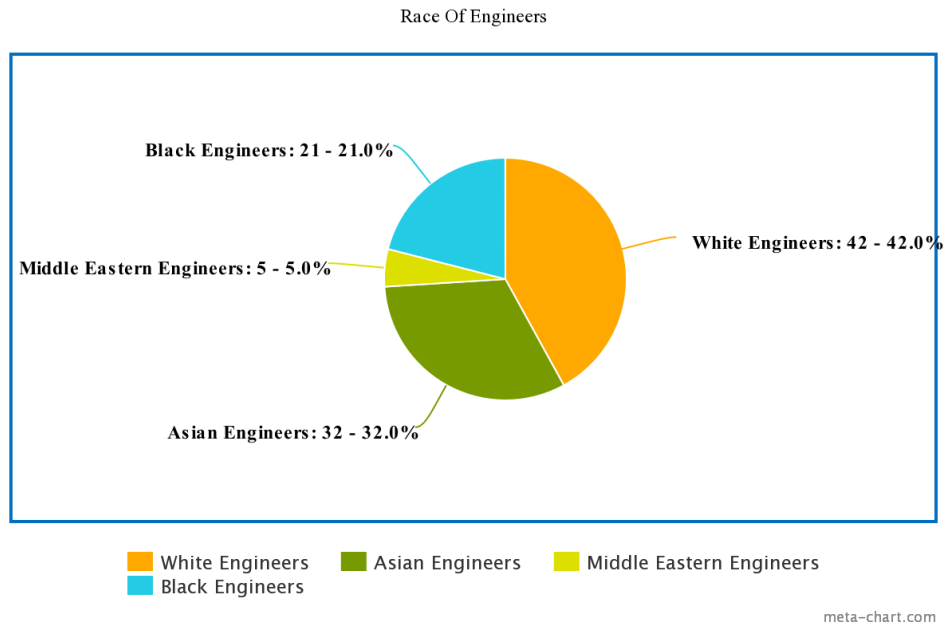
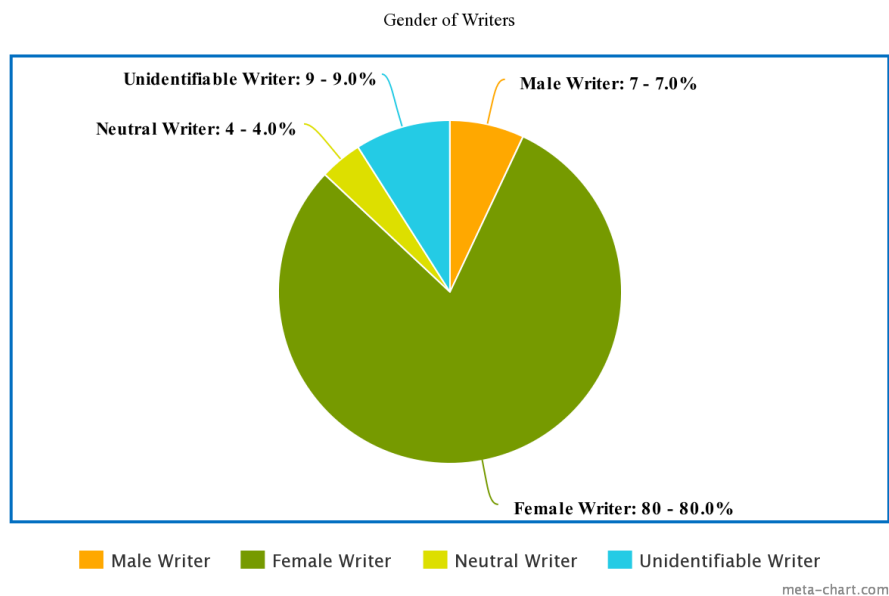


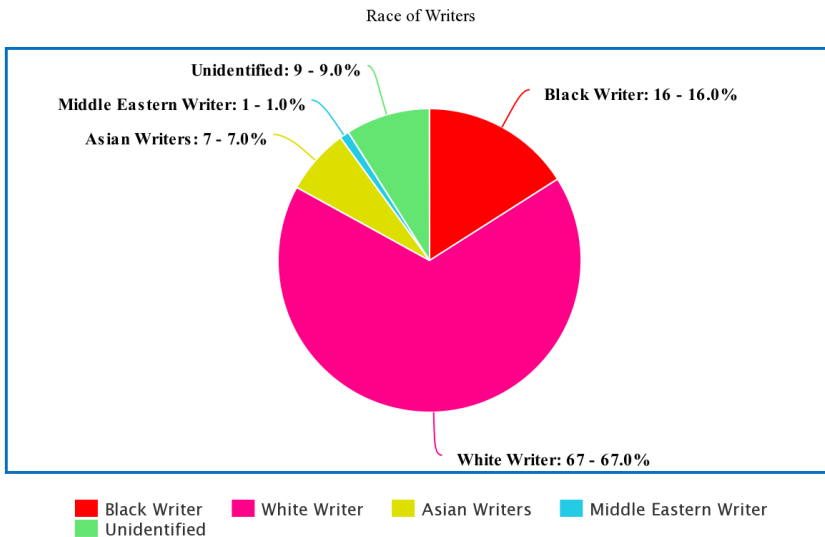
Figure 3. Race of Engineers

For the writers, Craiyon v3 generated a lot of abstract images, unlike Bing, which created mostly basic images of people on the job. The results ended with 80% of the writers being females, 7% being male, 4% being neutral, and 9% not following the prompt.

Figure 4. Gender of Writers



For the writers' races, most writers were depicted as white with 67%. It was followed by black with 16%, Asian with 7%, and Middle Eastern with 1%. The other 9% were outliers.



meta-chart.com Figure 5. Race of Writers

As for the age results, most of the writers depicted were between 18 and 30 years old as they made up 73 of the 100 images. The people generated that resembled to be between 30 and 60 took up 20% of the results. The other 7% were not following the prompt.

After generating 100 images of firefighters from Microsoft Bing image creator and analyzing the photos, We were able to break down the gender, race, and age. For the genders, 20 of the images had male firefighters, and 13 were female firefighters. The other 67 images were neutral as they couldn't be identified because of masks over their face.




Figure 6. Gender of Firefighters

Out of the 33 firefighter images that could be recognized, 20 of them were white, 3 were Asian and 7 of them were black.

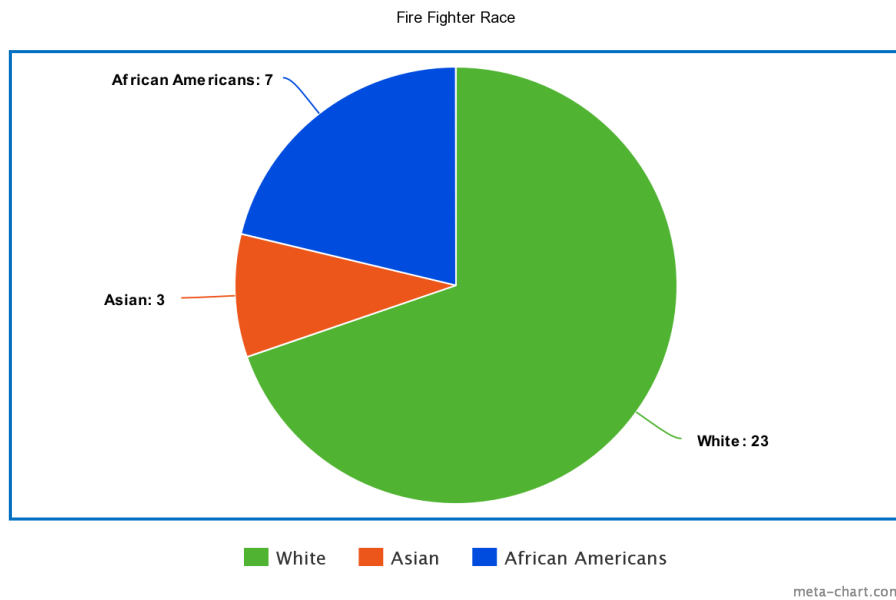


Figure 7. Race of Firefighter

For age, 32 out of the 33 of the firefighters were depicted as ages between 18 and 30. One image that was generated shows a male firefighter who appears to be older than 30.

Discussion

After reviewing all of the images, it is made abundantly clear that there is a race, gender, and age bias. Across the different careers, each showed one or multiple biases. Within the engineering category, it was recorded that the majority of the AI-generated images were female, and they were mostly within the 18-30 range when it came to age. It was also shown that the majority of the images showed people of white, or Asian race. According to Zippia “There are over 228,911 engineers currently employed in the United States. 13.7% of all engineers are women, while 86.3% are men. The average engineer age is 40 years old. The most common ethnicity of engineers is White (67.9%), followed by Asian (15.0%), Hispanic or Latino (9.1%) and Unknown (4.6%).” (*Engineer Demographics and Statistics [2023]: Number of Engineers in the US, 2023*).

The gender, and age data collected from Bing is the opposite of the actual statistics collected by Zippia. According to Bing engineers are mostly gorgeous young women. The male images that Bing created were mostly above the age of 30 and were handsome men. The race data collected is for the most part accurate to what Zippia stated, however, there is no Hispanic representation from AI-generated images.

Lastly, for firefighters, Zippia states “There are over 138,819 firemen currently employed in the United States. 11.8% of all firemen are women, while 88.2% are men. The average fireman age is 38 years old. The most common ethnicity of firemen is White (69.0%), followed by Hispanic or Latino (15.3%), Black or African American (8.1%) and Unknown (4.8%)”

(Fireman Demographics and Statistics [2023]: Number of Firemans in the US, 2023). There was not much of a gender bias in Bing as most of the images were neutral, and the identifiable ones were similar to actual statistics. Although there was age bias as most of the images generated appeared to be somewhere between 18 and 30 while Zippia states otherwise about the average firefighter age. The images created by Bing portray some of the male firefighters as muscular men with a burning house behind them. It also portrayed some of the female firefighters as attractive females saving cats. Bing however did display a race bias as most of the identifiable images were of white people, and once again there was no Hispanic representation within the AI-generated images.

For writers, Zippia reports that “There are over 45,024 writers currently employed in the United States. 53.8% of all writers are women, while 46.2% are men. The average writer's age is 42 years old. The most common ethnicity of writers is White (75.1%), followed by Hispanic or Latino (7.8%), Black or African American (6.1%) and Unknown (5.6%)” *(Writer Demographics and Statistics [2023]: Number of Writers in the US, 2023)*. From the data that was collected from Craiyon v3, it can be said that there is gender, race, and age bias. The AI software shows a gender bias due to most of the images which were received are female even though writers have an almost equal ratio between male to female, while we received many more women images compared to men. Craiyon v3 also displayed an age bias as “Zippia” claims that the average age of writers is 42 years old, while the AI images depicted most writers as younger adults. Craiyon v3 also displayed race bias as it generated a more equal number of people of different races, but did not show Hispanic representation just as the other prompts did.

While creating the images, it should be restated that most of the images were gathered from Bing. Due to this, it can't be confirmed whether there truly is bias within AI software, but

instead, it can be confirmed that there is bias within generated images from Bing. Despite this, the data which was gathered does not disprove the hypothesis.

Another paper that also discusses this problem is “Detecting Race and Gender Bias in visual representation of AI on web search engines”. In this paper, they discuss race and gender bias within AI web searches. As asserted by “Detecting Race and Gender Bias in Visual Representation of AI on Web Search Engines” it states “Our analysis showed that non-racialized portrayals of AI are prevalent on Western search engines (Fig.2). With the exception of Bing, where racialized images prevail among the first 10 results...Just as in the case of race, the majority of Western search engines (except Bing) prioritized gender-neutral images among the top 10 results with more gendered images appearing for lower results” (Makhortykh et al., 2021). “Detecting Race and Gender Bias in visual representation of AI on web search engines” supports the data which is present in this paper. It shows that sites such as Bing have a bias in both gender and race.

The data they collected was different from the data that we collected due to them using a different method. They used AI based on search methods. According to “Detecting Race and Gender Bias in Visual Representation of AI on Web Search Engines” “To collect data, we utilized a set of virtual agents - that is software simulating user browsing behavior (e.g., scrolling web pages and entering queries) and recording its outputs” (Makhortykh et al., 2021). Their data might have differed from our data due to them using search engines while we used text-to-image AI generation.

Conclusion

To conclude, technology is constantly evolving. But humans are the ones making these breakthroughs and evolutions. Due to this, gender, race, age biases, and stereotypes should leak

into the algorithms used to create AI. This should be made prevalent in text-to-image AI generation. After reviewing all the images generated from the two websites, it is apparent that there is gender, race, and age bias. Overall, the AI image software helped prove that there is bias within the algorithm. Moving forward, in future research around this topic, there should be more diversity among the different AI image generators. This would be beneficial to truly see if there is bias within AI, and help identify the problem.

References

Engineer Demographics and Statistics [2023]: Number Of Engineers In The US. (2023, July 21).

<https://www.zippia.com/engineer-jobs/demographics/>

Fireman Demographics and Statistics [2023]: Number Of Firemans In The US. (2023, July 21).

<https://www.zippia.com/fireman-jobs/demographics/>

Makhortykh, M., Urman, A., & Ulloa, R. (2021, April). Detecting race and gender bias in visual representation of AI on web search engines. In *International Workshop on Algorithmic Bias in Search and Recommendation* (pp. 36-50). Cham: Springer International Publishing.

Stypinska, J. (2023). AI ageism: a critical roadmap for studying age discrimination and exclusion in digitalized societies. *AI & society*, 38(2), 665-677.

Writer Demographics and Statistics [2023]: Number Of Writers In The US. (2023, July 21).

<https://www.zippia.com/writer-jobs/demographics/>