AI Image Generation Guilty of Perpetuating Racial & Gender Stereotypes

Aashir Ali, Sehajdeep Singh, Sajid Rahman

Department of English, City College of New York

ENGL 21007: Writing for Engineering

Pamela Stemberg

April 4th, 2025

Abstract

This study investigates gender and racial biases in AI-generated profession imagery using Craiyon AI. The experiment generated 300 images of three professions (doctor, lawyer, and college professor) and compared the demographics depicted in the images to real-world data from Zippia.com. The findings revealed a strong gender bias, with males overrepresented in typically male-dominated roles. The AI showed limited diversity as it skewed away from African Americans in all of the professions and instead showed a heavy bias towards Whites. Nonetheless, the results of this study are consistent with those done in the past, and our study proves that there is indeed a bias in AI image generation.

Introduction

Artificial intelligence emerged as a storm in this world. It is a compelling tool that can be used for various applications and allows users to create many things in the span of a few seconds. This lab will primarily focus on the image creation aspect of the AI world and mainly focus on the underlying bias in AI. There are many platforms that allow users to create images, such as Dalle, Crayon, Bing, and Canva's AI image generator. Even though these tools allow endless possibilities, growing evidence shows that these systems often reflect and reinforce racial, gender, and age bias. When asking these AIs to construct images on certain topics like different professions/jobs, AI-generated images frequently show concerning results that highlight racial/sexist stereotypes in certain situations and jobs.

Some recent studies that highlighted these troubling patterns include the article, "Racial bias in AI-generated images" by Yiran Yang who conducted a large-scale study that analyzed a Chinese AI-powered image generator called Meitu. The findings convey that the platform

depicted white individuals a lot more accurately than people of color. This was the case in many different racial contexts/situations, which just goes to show the AI pushing this agenda of "white normality" as it had the highest accuracy to depict images as white people. It goes on to reflect the deeper embedded issue in these image generation algorithms that are provided by their parent/development companies.

Similarly, the study, "Bias in Generative AI" conducted by Zhour et al. had analyzed many AI-generated images from platforms like Midjourney, DALL-E 2, etc. When asking it for roles that relate to being "successful", the platform mostly produced young white male figures to represent these prompts. This reflects not only a bias but also a huge discrepancy in the training data of these models and not having proper diversity data. Women and darker-skinned people were significantly underrepresented in many scenarios; for example, when asking the AI to draw portrayals of authority figures like CEOs or scientists, most of the time it generated white males. It also depicted women as young and happy while depicting males as old and grim and while this is not as pertinent as the previous example, it still can show the AI convey women as submissive and less complement than their male counterparts. (Zhour et al, 2024)

Another study by Bloomberg further discusses this issue in AI image generation. Bloomberg's reporters discovered that inputting prompts like "CEO" and "lawyer" resulted in white men, while prompts like "nurse" or "teacher" showed majority white females. Even the attempts made to specify the race in the prompts led to inaccurate results, suggesting the AI's inability to represent gender and race accurately in roles that are gender neutral. (Humans Are Biased… , 2023)

These findings altogether show a pattern of bias in these AIs, which can go on to cause severe consequences if left unchecked. Due to the booming popularity of AI in today's age, it can

shape public perceptions and reinforce racial inequalities and overall cause bigger issues rather than solving them.

Hypothesis

We hypothesize that AI image generators will primarily construct certain stereotypes of different roles. For example, roles like lawyers and doctors will primarily show white males, while roles like professor and teacher will primarily be female-predominant. To test this, we will research using the AI image generator Craiyon and research gender-neutral professions to see what the AI constructs. This should clearly display the bias AI has on different demographics and shed light on this issue.

Materials and Methods

Materials

- Computer
- Wifi (Access to the internet)
- Microsoft Excel
- Craiyon (image generator)
- Google Drive

Method

For the image generation platform, the team decided to use Craiyon throughout the entire experiment. The study's results, which will be discussed in the following section, were obtained using the following procedure. First, the team created a Google Drive and made three separate folders within the drive and named them “Doctor,” “Lawyer,” and “College Professors” to be organized. After that, each member of the team created a unique folder on their computers

regarding the profession they were doing. For example, the person generating images for lawyers created a folder on their computer named “Lawyers” and used it to save images into that folder when done generating. Next, we went on Craiyon, where we started the image generation process. To generate images of college professors, the team used the term “college professors” and used the v4 Auto style since it was the default. When the AI finished generating the images, we clicked on each image and then “save as” and renamed each image according to the gender and race of the person in the picture. For example, “white male (1).” The 1 is there as a way to measure how many white males we have in the current profession. The team later used this as a reference when creating pie charts and tables. We repeated these steps 12 times for each profession to generate 100 images for each category, which resulted in 300 total images.

Once the images were all saved in each folder on our computers, we then went back on Google Drive and dragged and dropped all three folders into the drive. The drive did its work and saved all 300 images there. Now, the team had 3 separate folders, “Doctor,” “Lawyer,” and “College Professors,” with 100 images in each folder, with the name of each image corresponding to the gender and race of that specific person in the photo. Once this part was done, we then shifted our focus to Excel sheets. We created 6 sheets in one Excel book that followed the pattern of “Name of the profession, gender” and “Name of the profession, race.” So for example, “Doctors, gender” and “Doctors, race.” We used the data from our Google Drive and first created a table for college professors’ genders. We named the first cell A1 “Gender” and the adjacent cell B1 “# of people.” Then, we named the cells A2, A3, and A4 “Male,” “Female,” and “Other,” respectively. After that, we manually wrote our data into the adjacent columns for each gender. Then, we went to “insert” in the top left corner and clicked on “recommended

charts" and chose the 3rd recommended option, which was a sideways bar graph. Once the bar graph was displayed, we renamed the title of the graph "Genders of College Professors." The team used this as the main way to showcase data for the genders for each profession.

After this, we shifted our focus onto the data we obtained for race. We went onto another Excel sheet in the same book, but this time for "College Professor's Race". We named the first cell A1 "Race" and the adjacent cell B1 "# of people". Then, we named the cells A2, A3, A4, and A5 "Asian", "White", "Black", and "Hispanic" respectively. Once this table was set up, the team manually wrote in the data for each column by referring back to the Google Drive. After finishing that, we clicked on "insert" and "recommended charts" again, but this time we chose the first option which was a pie chart. After the pie chart was displayed, we clicked on the plus symbol next to it, then the data label, and then more options, and checked off "display percentages". Additionally, we renamed the title of the chart to "Race Of College Professors." The team repeated all of these steps twice more for the remaining professions (doctors and lawyers).

Results

This part of the results section presents the gender and race distributions that we obtained from our experiment of professions across three fields — Lawyers, College Professors, and Doctors — in graphical formats. These visual aids help interpret patterns and disparities across gender identities in each occupation.

Figure 1

Gender Distribution of Lawyers

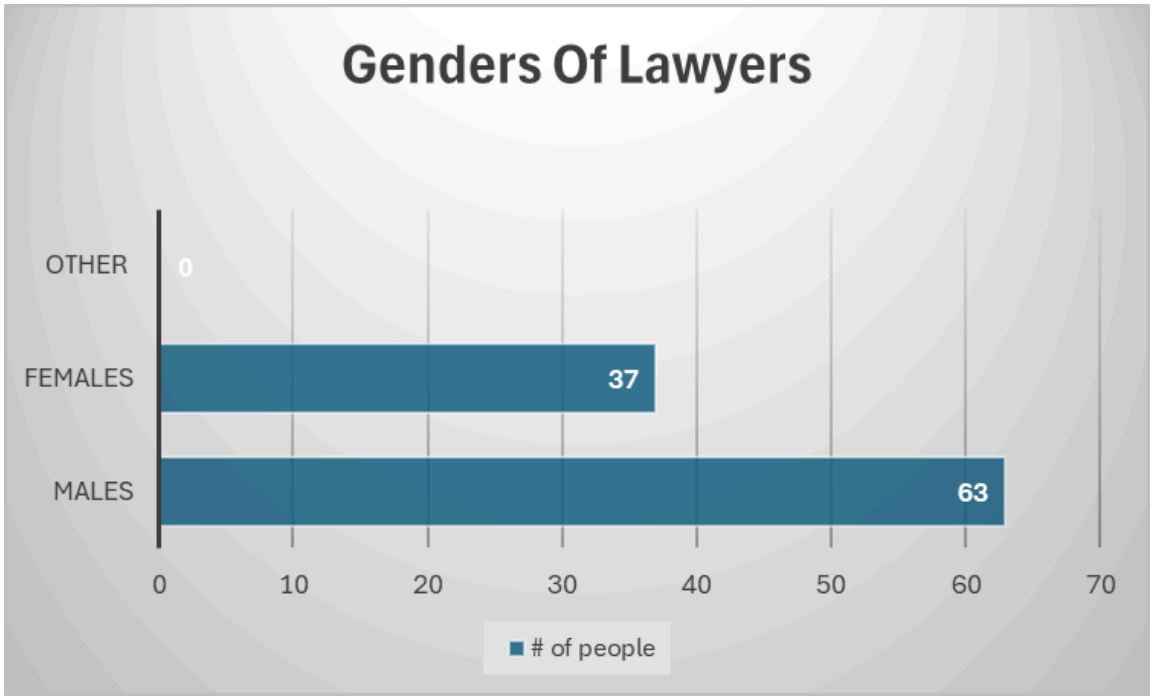

Note. Number of people per gender occupying the Lawyer industry.

Figure 2

Gender Distribution of College Professors

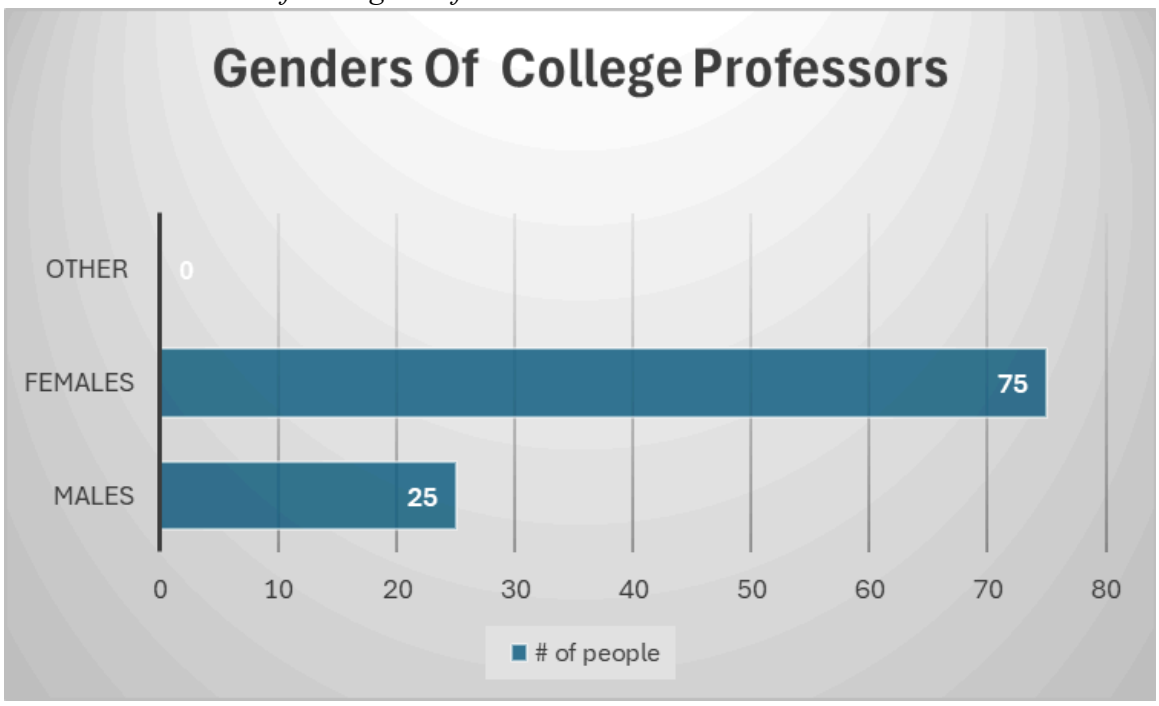

Note. Number of people per gender occupying the College Professor industry.

Figure 3

Gender Distribution of Doctors

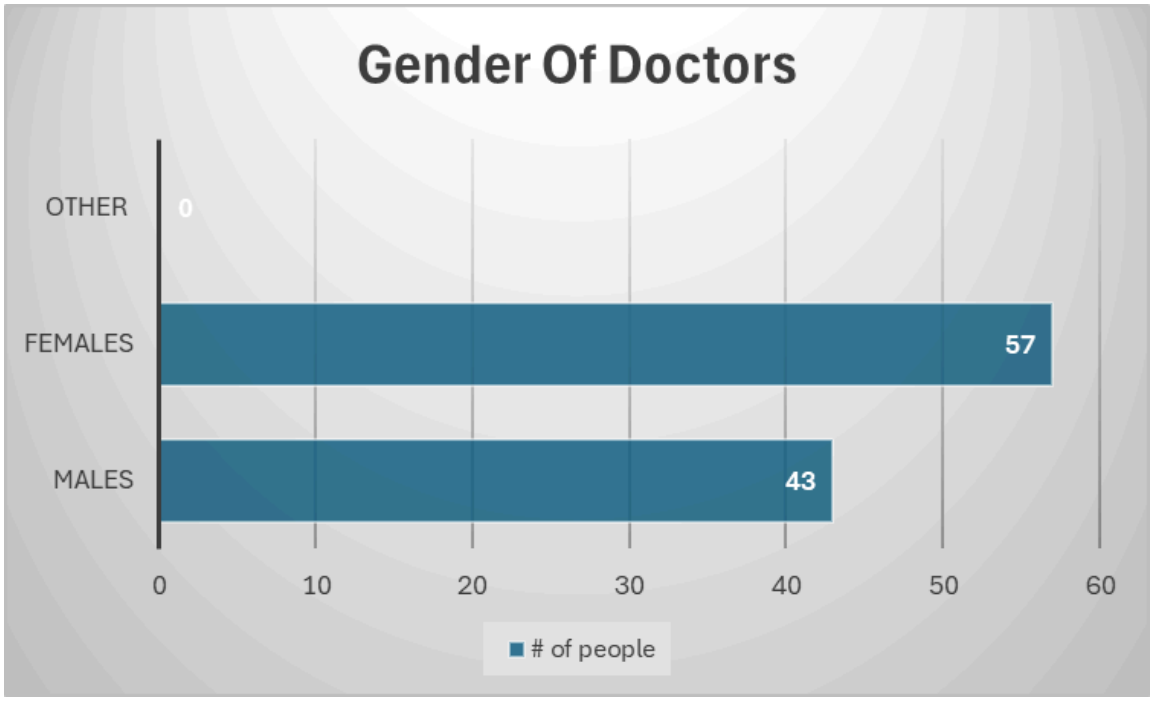

Note. Number of people per gender occupying the Doctor industry.

Figure 4

Race Distribution of Lawyers

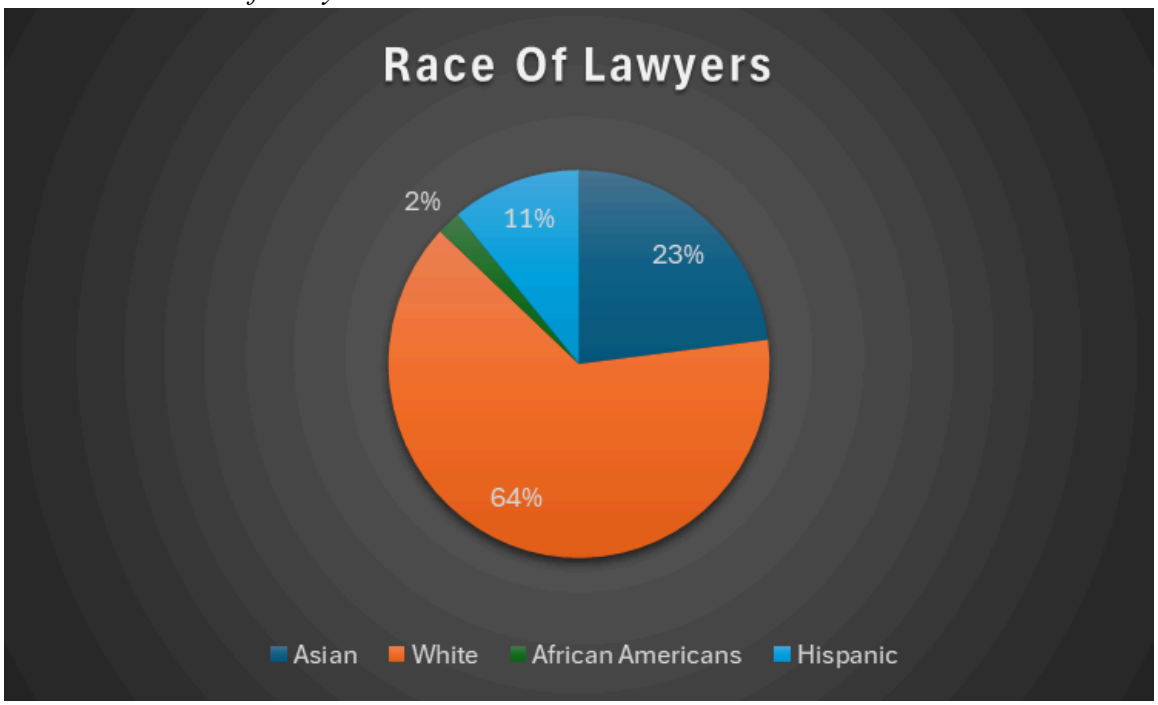

Note. Percentages per race that occupies the Lawyer industry.

Figure 5

Race Distribution of College Professors

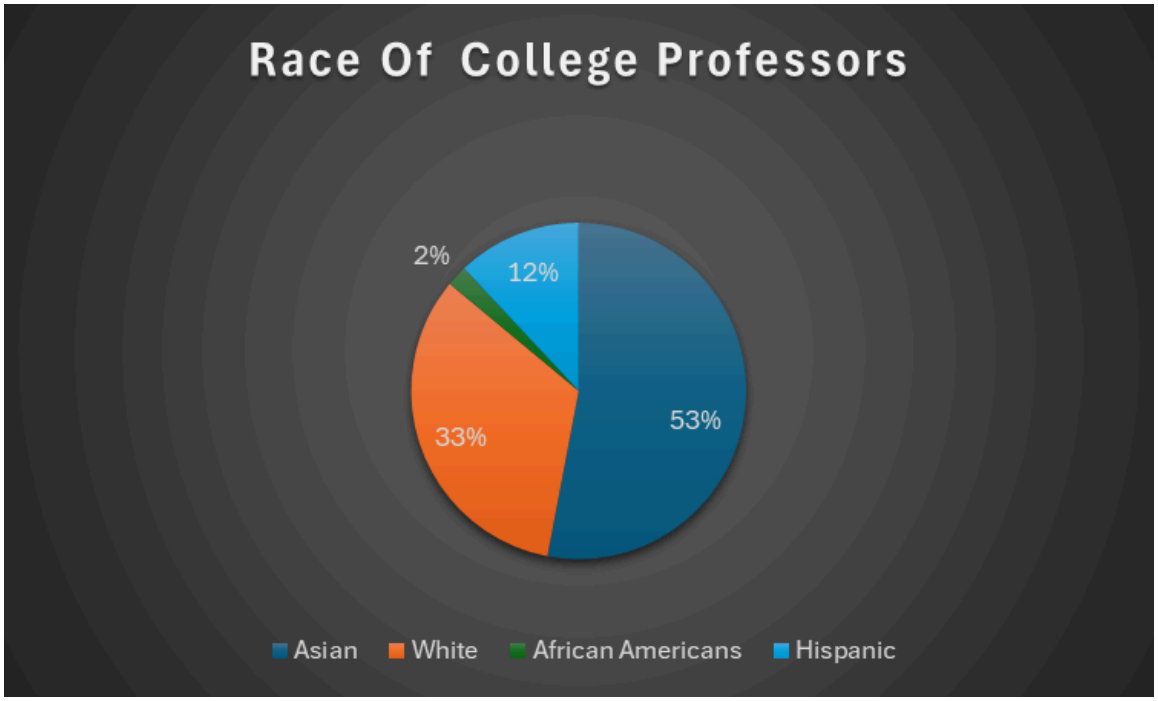

Note. Percentages per race that occupies the College Professor industry.

Figure 6

Race Distribution of Doctors

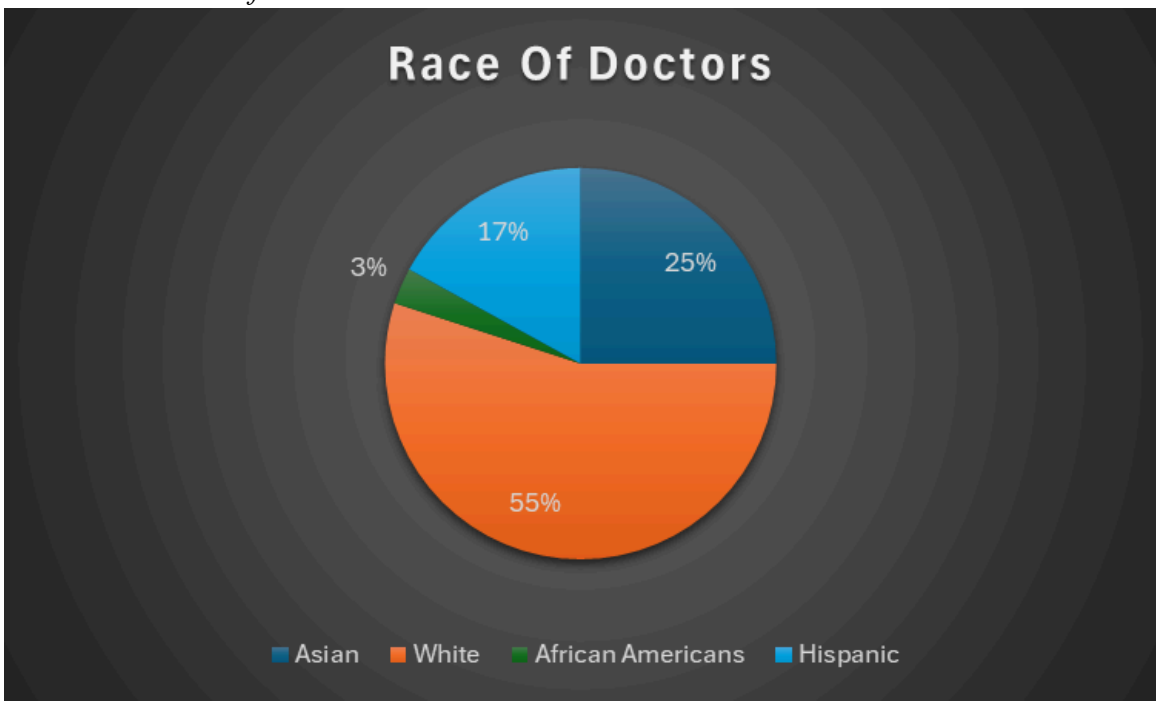

Note. Percentages per race that occupies the Doctor industry.

This part of the results section presents the gender and race distributions that we obtained from **Zippia** for the professions — Lawyers, College Professors, and Doctors — in graphical formats. These visual aids help interpret similarities and differences between our own data.

Figure 7

Gender Distribution of Lawyers from Zippia

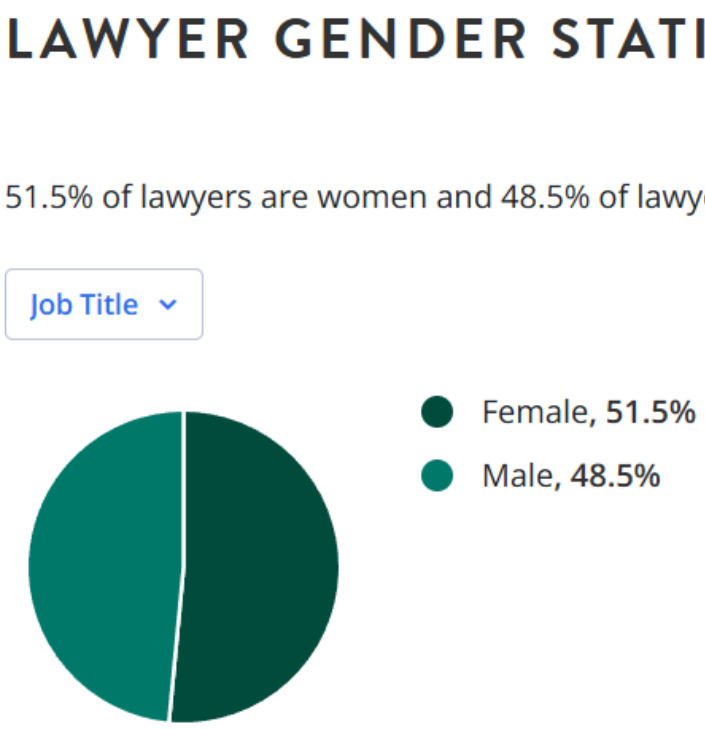

Note. Percentage breakdown of genders occupying the Lawyer industry. From Zippia (2025)

Figure 8

Gender Distribution of College Professors from Zippia

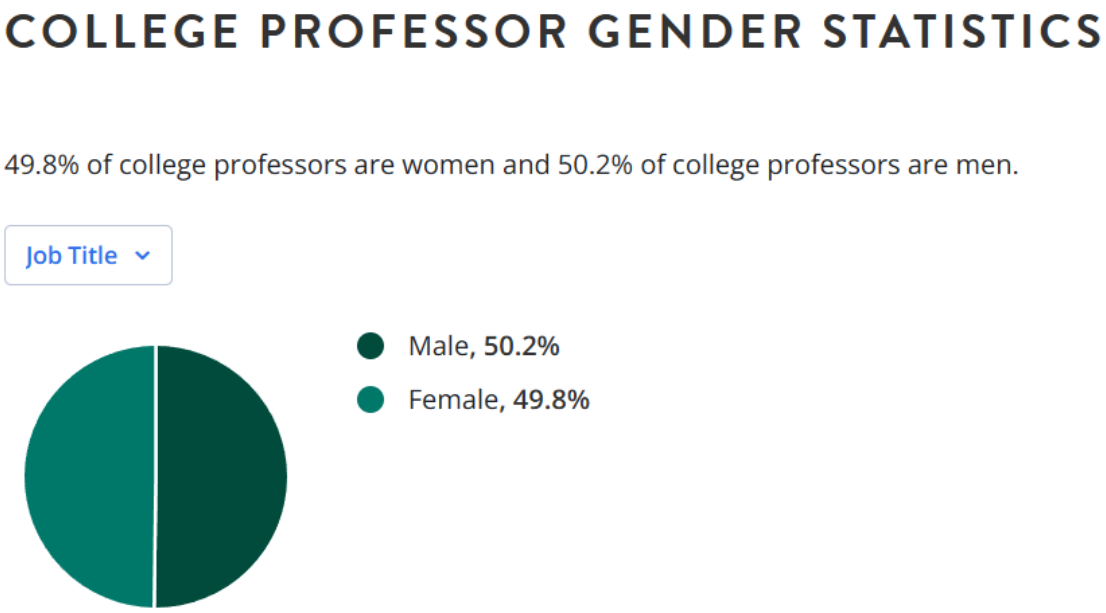

Note. Percentage breakdown of genders occupying the College Professor industry. From Zippia (2025)

Figure 9

Gender Distribution of Doctors from Zippia

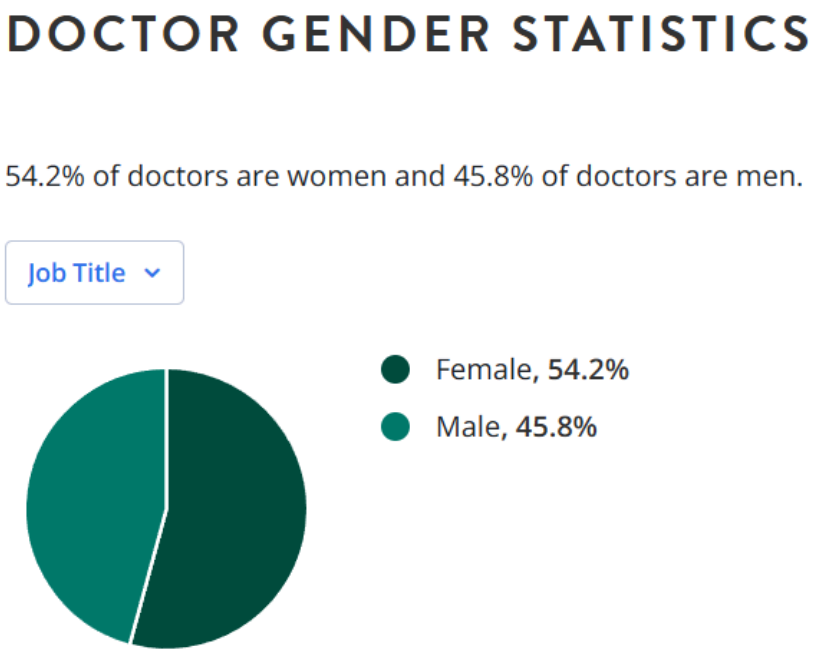

Note. Percentage breakdown of genders occupying the Doctor industry. From Zippia (2025)

Figure 10

Race Distribution of Lawyers

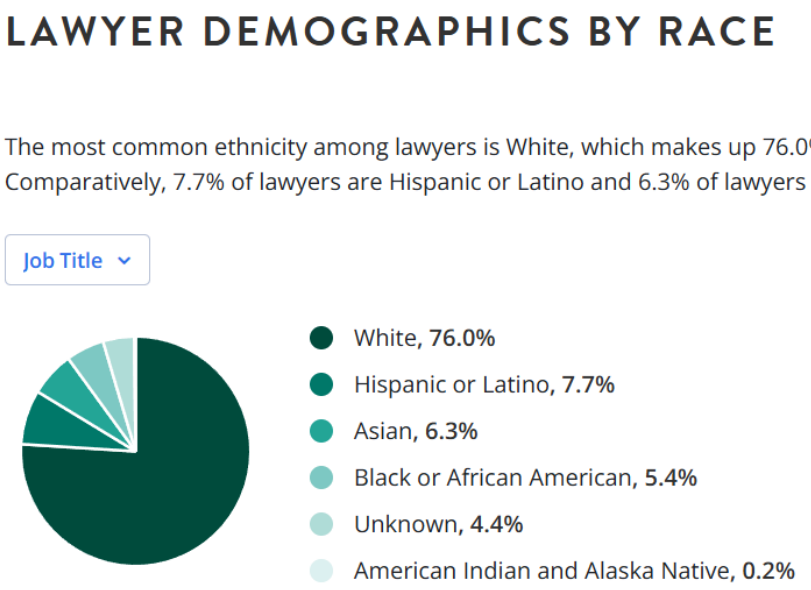

Note. Percentage breakdown of race occupying the Lawyer industry. From Zippia (2025)

Figure 11

Race Distribution of College Professors

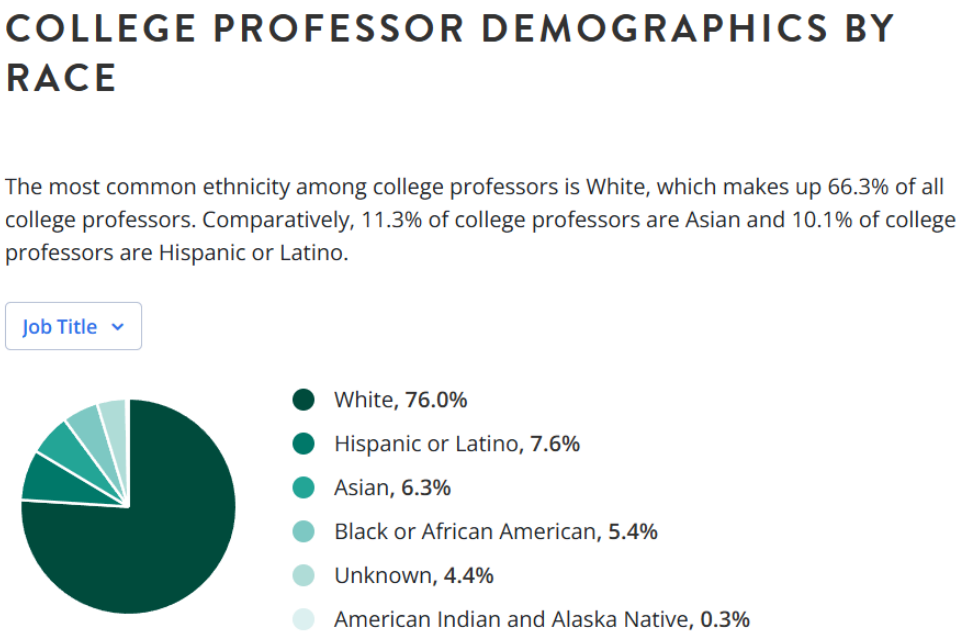

Note. Percentage breakdown of genders occupying the College Professor industry. From Zippia (2025)

Figure 12

Race Distribution of Doctors

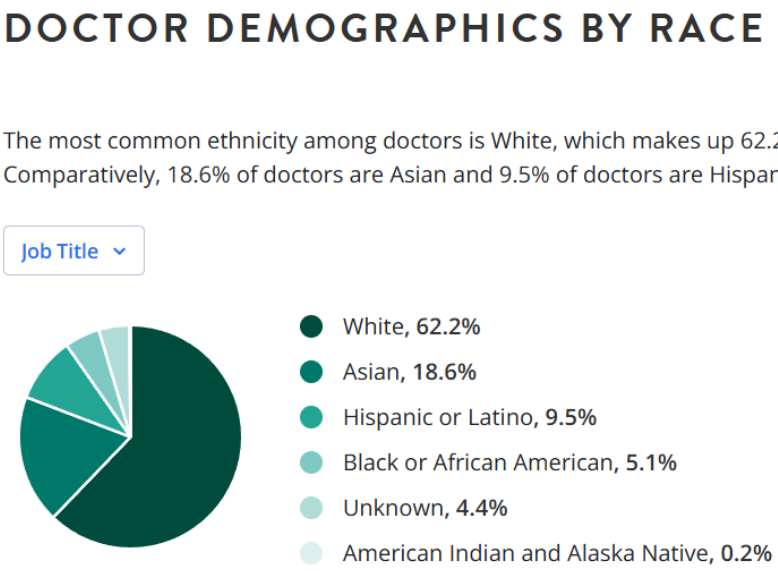

Note. Percentage breakdown of genders occupying the Doctor industry. From Zippia (2025)

Discussion/Conclusion

The results obtained from our experiment clearly indicated a sense of stereotype and bias. In our image generation for college professors, from Figure 2, we attained 75 females compared to only 25 males. That's a 75% to 25% comparison in terms of percentages. From Zippia, in Figure 8, we see that in reality, there are about 50.2% females in the industry compared to 49.8% males (College Professors demographics and statistics in the US, 2025). The AI clearly shows a bias and stereotypes the college professor's role towards women. This can be due to several reasons. One of the reasons may be due to a sense of comfort that people get when females teach compared to men. Females tend to exhibit more emotional intelligence and therefore can connect with students on a deeper level. As for race, in Figure 5 our experiment shows that 53% of images generated were Asian, 2% were Black, 12% were Hispanic, and 33% were White. From Zippia, in Figure 11, 76% of professors are white, 7.6% are Hispanic, 63% are Asian, and 5.4%

are African American (College Professors demographics and statistics in the US, 2025). This shows a drastic shift from what we got from our experiment with the AI to those that Zippia got from reality. The AI showed bias towards Asians for this profession, clearly. The reason for this could be because AIs are trained on vast amounts of information based on real-world statistics, and it's true that Asians are usually on top of the leaderboard when it comes to intelligence and jobs, so the AI used that information to generate images for professors, who are believed to be some of the smartest people in their industry.

Moving onto lawyers, in Figure 1 in our experiment, we saw that 63 out of 100 images the AI generated were all males. Only 37% were females. According to Zippia, in Figure 7, 51.5% of lawyers are women, while 48.5% are men (Lawyer demographics and statistics in the US, 2025). Our AI clearly showed a bias towards males in this case. Furthermore, in Figure 4, the AI generator generated 64% whites, 23% Asians, 11% Hispanics, and only 2% blacks. Now, compared to the findings of Zippia, in Figure 10, 76% of lawyers are white, 7.6% are of Hispanic descent, 6.3% are Asian, and 5.4% are African American, while the rest were labeled as unknown (Lawyer demographics and statistics in the US, 2025). There is a notable similarity between our data and the data from Zippia for race. Our data was only 12% off for the amount of white lawyers. Additionally, we were about 3.4% off for Hispanic lawyers. This shows that the AI generator was strong in regard to the lawyer profession.

Moving onto the final profession, doctors, in Figure 3, we can see that the AI exhibited some bias towards females as it generated 57% of them compared to 43% males. Now compare this to the findings of Zippia in Figure 9: 54.2% were labeled as females and 48.5% as males. This shows that in reality, the AI wasn't that off and didn't really stereotype any gender, as it accurately generated enough images for both genders. Transitioning to race, in Figure 6, we can

see that the AI once again showed bias towards whites as it generated 55% of images for that race. The rest were 25% Asians, 17% Hispanics, and only 3% Black. Comparing this to Zippia, in Figure 12, 62.2% of all doctors were found to be white, 18.6% were Asian, 9.5% were Hispanic, 5.1% were African American, and the rest were deemed unknown (Doctor demographics and statistics in the US, 2025). Here we can see that the AI, in fact, was very close in terms of race. The numbers aren't that off and this could be due to the source that it's feeding information off of. A key observation that we noticed is that the AI didn't show any favoritism towards African Americans in any professions. This could be due to a vast amount of reasons, but one reason that particularly comes to mind is that in the real world these professions that we've chosen don't favor Africans in terms of job demographics. We saw from Zippia's data that even they had African Americans with the lowest percentages. As far as issues, there were some issues with the AI generation platform, Craiyon. At times it generated something completely off-topic. Other than that, there were no other real issues and we're satisfied with our study.

AI has come a long way. From once being a myth and something far in the distance, it has been brought to life by the hard work of many platforms and it's able to do some incredible work. However, it still has a long way to go and there is always room for improvement.

Self-Reflection

Aashir Ali

When doing the experiment and seeing the results generated right in front of my eyes, I wouldn't say I was really shocked. I kind of expected AI to be "racist" in a sense because AI is trained on a vast amount of databases that are collected from real-life statistics. For example, for college professors, I already expected Craiyon to generate mostly Asians, and it in fact did. This

is because in leaderboards such as the ones about what race scored the highest on the SAT on average, you'll see Asians right around the top 3 if not number 1. Now AIs look at that information and correspond that type of information to college professors because they've been trained to think that college professors are really smart, which they are, and by connecting the dots, the AI is able to associate races that rank high in test scoring with college professors.

Sehajdeep Singh

This lab report really helped me to understand how AI is developing in our world and what is actually going on in the background. My whole perspective with AI was that it is a really helpful tool that is great for doing and learning many things. But after conducting the research and producing the conclusion that we made in this lab, it really helped to open my eyes as to what is actually going on. I knew that when it came to AI image generation, there was going to be a sort of bias or racist ideal being shown when portraying people, but the outcomes that we achieved really shocked me due to how apparent it actually was. When also looking at research on the same topic from different primary sources, the outcome was similar, as all the research showed the AI pushing stereotypes to the user even when being told which race to draw.

But overall I really liked the whole lab report experience, as it was great to delve into these types of writing skills and to actually be able to produce a professional lab report on a topic. We all had to work together, combine our research, and reach a conclusion by comparing it to actual data sets. The main thing I can take away from this lab report is the fact that many minorities are underrepresented in AI image creation, which makes you think about what type of training data these AIs were being developed on and why. A lot of the more successful professions have white males being portrayed in them, while other professions, like social

worker or janitor, have darker skin tones being produced. This similar trend of stereotyping is seen in other scenarios, like when asking the AI to draw a group of people and it being mostly white or even the fact that it needed a command in its code to make sure it diversified something when that really shouldn't be needed. Overall it is definitely interesting to see these outcomes, and they definitely need to be brought to light so it doesn't get worse and doesn't spread these ideals to people even more.

Sajid Rahman

My understanding of artificial intelligence was tested while working on the AI Image Generation Lab Report, which also made me more conscious of the societal ramifications of emerging technologies. Our project, which looked at racial and gender biases in AI-generated images of three professions, made me consider how data, training models, and algorithms can inadvertently reinforce negative stereotypes.

Working with my team was one of the most fulfilling parts of this assignment. With each member concentrating on a particular profession, we divided the duties. We were able to contribute to a common objective while working independently thanks to this structure. I learned how important communication and organization are when working on group projects. Every stage, from organizing folders to classifying picture data according to gender and race, called for accuracy and collaboration. I also used programs like Excel to create visual data, particularly in the areas of organizing, analyzing, and visualizing statistical data.

Keeping objectivity when classifying AI-generated images was the most difficult aspect. There were times when Craiyon's output was unclear or failed to explicitly represent a particular race or gender. It also generated something completely irrelevant and not related to the profession. This encounter has caused me to reevaluate how our society measures bias and

representation, as well as how AI systems frequently uphold prevailing narratives while posing as neutral.

It was especially interesting to compare our results with actual data from Zippia.com. It became clear that even though AI has advanced remarkably, issues with representation and fairness still exist. African Americans, for example, were continuously underrepresented in all of the professions we examined, reflecting a larger systemic problem. The lack of corrective nuance in AI generation raises concerns about how these tools could affect public perception if left unchecked, even though this may reflect real-world statistics.

References

College Professor Demographics and Statistics [2025]: Number Of College Professors In The US. (2021, January 29). Zippia.com. https://www.zippia.com/college-professor-jobs/demographics/#gender-statistics

Nicoletti, L., & Bass, D. (2023, June 9). Humans Are Biased. Generative AI Is Even Worse. *Bloomberg*. https://www.bloomberg.com/graphics/2023-generative-ai-bias/

Yang, Y. (2025). Racial bias in AI-generated images. *AI & SOCIETY*. https://doi.org/10.1007/s00146-025-02282-1

Zhou, M., Abhishek, V., Derdenger, T., Kim, J., & Srinivasan, K. (2024). Bias in Generative AI. *ArXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2403.02726

Zippia. (2021a, January 29). *Doctor Demographics and Statistics [2021]: Number Of Doctors In The US*. Www.zippia.com. https://www.zippia.com/doctor-jobs/demographics/

Zippia. (2021b, January 29). *Lawyer Demographics and Statistics [2021]: Number Of Lawyers In The US*. Www.zippia.com. https://www.zippia.com/lawyer-jobs/demographics/