AI Image Generation Study Assignment

Authors: Jayden Mallari, Josue Rodriguez, David Mayancela Suquila
Date: 7/15/2026

Table of Contents

Abstract
AI may potentially have biases based on the data it has been trained on, causing it to generate images or other results that are skewed in their appearance for or against certain demographics. Using Gemini AI's Nanobanana image generation model, and prompting it to create 100 images of three different professions, we will focus on Political-Senators, Journalists, and Public Defenders. It was found that these biases do exist to different intensities. The model has a tendency to misrepresent or overrepresent certain demographics, often based on historical precedence. Nominally, Male White-Europeans, Female White-Europeans, and Female African-Americans were found to be the most commonly generated of the three respective professions. Moreover, the assumptions made by the AI model tend to be accurate when it comes to the age range of the professionals in each category. Evidently, when the generative model tends to portray the persons of each profession based on the North-American standards of beauty. This is because bias could appear as a hidden lawyer,it is important to notice that "thought shortcuts" the generative ai is mimicking from the data which it was trained on.

Introduction:

Today, anyone can become practically anything they desire when it comes to the job market. Previously positioned gender roles that lasted only as far as the later end of the last century have become redundant as both education and tolerance for individuals has become significantly more loose. Not only that, but ethnicity and age also no longer presents a barrier to higher jobs that were previously only attainable by certain individuals. Despite that, there are many holdover norms from those days that tend to cause those professions to only be viewed as being for the demographic they were previously most heavily populated by. This isn't just perceived by humans, but also likely by AI and algorithms as well. Since its major boom in 2023, it has been widely used in many settings, between standard discussion, problem-solving, and even replacement of whole departments or assets in companies for more efficient workflow. It has been trained on an intensive number of datasets over the years, which permits it to do many of the functions it is capable of now. Nevertheless, historical knowledge has let us understand that there are prejudices, because datasets contain historical information in the form of data then Datasets contain its patterns and trends too (CrashCourse, 2019, 1:15) . But even without human jerryrigging or otherwise, its generations have some internal biases and predispositions that most likely were of its own creation. The objective of the paper is to make an investigation about the accuracy of generative AI (text-to-image), in this way we could observe if there exist underrepresentation or misrepresentation of demographic groups in professional jobs positions. The hypothesis in the paper is that our data will show an overrepresentation of male in professional positions.

When we look for an image in google. The research engine will show us all the images related to the topic we are looking for. In this way the same word could have more than one meaning depending on the geographic and historical context. .

Before the models that could generate images based on text were especially trained to represent one type of image. Nonetheless, large Text to Image models are trained in large scale datasets that are often imbalanced and lacking demographic diversity, in consequence AI-generated imagery has the potential to reinforce stereotypical representations ((Mao et al., 2026). Because the outputs are based on the provided data. In this sense the LLM only knows the patterns associated with the words used. Because something being data is not information, in this sense the patterns found in the data could generate a tendency to incline the answers in a particular perspective. According to Mao one of the relevant implications of biased answers is the continuity of social perceptions that don't reflect the current world. Moreover as Mi Zhou mentioned in "Bias in Generative AI" : innovation continues its exponential growth it is important to ensure equity in the development and design of AI technology. Because AI is able to find patterns in large amounts of data there is another critical point when AI is asked about optimizing profits then is the implication of AI containing bias in itself.

An example of real-world implications is the reported in the Articles from WIRED (by Tom Simonite, 2019) and UChicago News (by Sam Jemielity, 2019) who corroborate the progressive incidence of privilege based on race being conducted by an AI algorithm intended to prioritize medical treatment in hospitals in the United States from as early as 2019. It was found that the system "effectively reduced the proportion of black patients receiving extra help by more than half, from almost 50 percent to less than 20 percent." (Simonite, para. 5). This was noted as being for possible financial reasons, as the algorithm prioritized the potential for being able to pay back treatment over definitive level of sickness, which would make it skew most towards white patients (Jemielity, para. 14.). It would seem that since its earliest times as simple systems of algorithms, AI was in some way susceptible to skewing data. Even with the time difference since the UChicago article's release, and even especially after the WIRED article.The Databases used contain the inclinations and biases from the people from which the data was collected.

Hypothesis

There is viable reason to believe that outside forces have an effect on how AI comes to its conclusions. Especially when it comes to image generation, where it may be more likely to be influenced to create something unusual. For this case in specific, the position taken is that AI is susceptible to influence based on the datasets it is fed, which may skew its subsequent generations. As a result, it is expected that some aspects of whatever it creates will not mirror real statistics.

Materials

For our case, we utilized the Gemini AI Nanobanana image generator, as it was the most advanced yet readily accessible model available to us. All documented data was to be placed on a Google Sheets spread, with rows denoting differences in age, gender, and ethnicity. The Sheets document, along with graphs and images generated by the AI were stored in a Google Drive for processing and finalization in this document.

Method

To provide a widespread view on the potential for biases within the AI, we picked three major professions with a known history of demographic stereotyping.

In order to analyze the bias

These are the positions of a Political Senator, a Journalist and a Public Defender. We then prompted the AI with the inquiry “Generate a photo-realistic image of a…”, with either of the three positions listed afterwards. At this point it would be important to note that between different models of AI, vastly different datasets have been used for training and prompt generation. As a result, the Gemini Nanobanana model may not hold the same biases as other models would.

After these generations, we looked at, and then discerned the images into different categories based on the face-level traits they possessed. During times of difficulty in telling what respective categories they would fall under, we would hold a consensus on the features of the images and then put them where they best fit. Although certain outputs followed a pattern in their appearance, the majority of them did have enough discernible features to warrant enough of these convenings for consensus, and despite that we did have enough information to imply different biases across each of the professions chosen for this study.

Results

Section A: Political Senator

Particularly we investigate the degree of bias in the generated images. In this case we explore the profession of politicians, particularly the US Senate, in which 100 senators participate (2 per State). These findings could contrast the historical precedent of 26% of senators were women and 74% of the senators were male (. As it is shown in the Woman Senators article from the United States Senate there have been 65 women in the senate and 1660 male senators, but once we look at the current year 2026 then we notice that there are 26 women and 74 men Senators in the 119th Congress (United States Senate, n.d.) . The results in Figure 1 show that there is a higher tendency in AI generation to generate the image of a female senator (77%) than a male senator (23%).

AGE

The images generated reflect a tendency to portray a female as senator even though they are a minority group in the Senate.

ETHNICITY

Based on Figure A2 and A3, we can discover that when it comes to the demographics distribution of Senate members the Generative AI (Nanobanana model) has a significant propensity to depict a US Senator as a White non-hispanic individual.

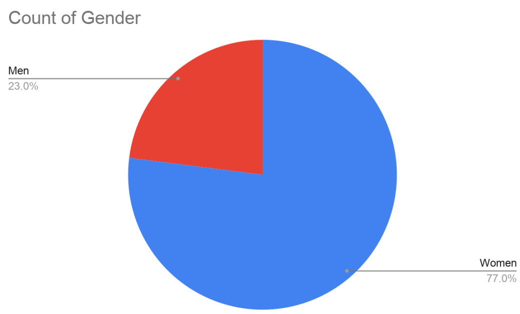

Figure A1: Gender representation

Additionally, another important item to note is that a near two-thirds majority of the senators are portrayed with a smile rather than a serious expression in the generated images, as depicted in figure A4. Further among the line of details in each image, the number of flags in the image tends to be a model value of 2, the maximum being 5 and the minimum zero (even in this case, there are still symbols and references to the US Senate).

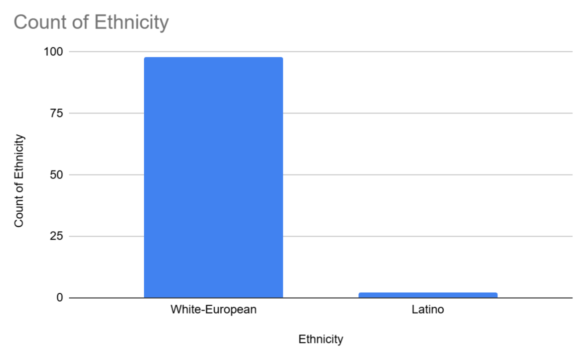

Figure A2: Ethnicity Distribution

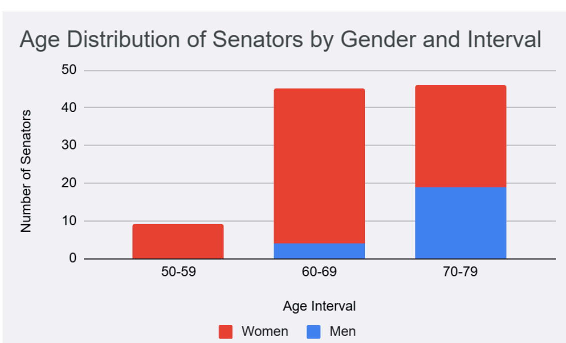

Figure A3: Age Distribution

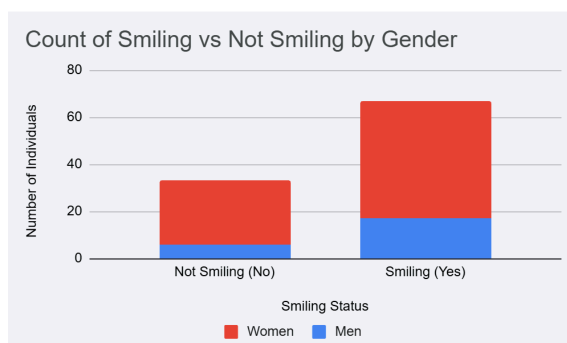

Figure A4: Count of Smiles

Discussion

AGE

The images generated reflect a tendency to portray a female as senator even though they are a minority group in the Senate.

ETHNICITY

Based on Figure A2 and A3, we can discover that when it comes to the demographics distribution of Senate members the Generative AI (Nanobanana model) has a significant propensity to depict a US Senator as a White non-hispanic individual.

According to the United States Senate (n.d.) the genders in the current 119th Congress is women and men.

AI did not reflect the diversity that there is in the Senate. Nonetheless, generative AI accurately describes the type of cloth and age that people in the Senate have. There is a tendency in generative AI to make Sanator standing behind a lectern and with a warm smile, the images appear to be happening while the Senator is addressing a topic to the public. Despite AI following the stereotype of the American politician, there is an opportunity for the database to be updated with data that reflects the diversity of the Senate in 2026.

Section B: Journalist

In this section of the study, the aim was to find potential biases in the way that AI generates, and likely to some extent perceives the average journalist, with no particular specification towards any region or individual. Despite that, the entire array of generations by Gemini Nano Banana were 100% female, with a majority of the images taking place in a city with protests in the background, or in a presumed warzone. The remaining minority take place in a nondescript European city. Additionally, aside from a few generations, the majority depicted women with features that would adhere to beauty standards.

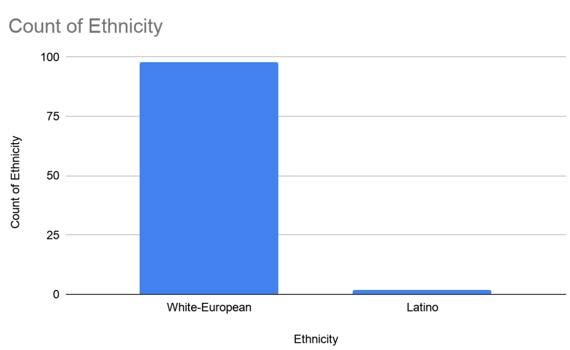

Figure B1: Ethnicity Distribution

Of all possible demographics within this study, only two were shown throughout the 100 images generated, with those groups being the White-European and Latino groups. Notably, the AI heavily skewed towards showing White-European denominations regardless of the setting, with an overwhelming 83% of individuals within the images appearing to be of that descent, and the remainder composing the Latino group (as depicted in Figure B1). This would be inaccurate to real-world demographics, as a 2023 study by the Pew Research Center indicates that overall, 51% of American journalists are male (Tomasik, et al., para. 6). Within that same study, however, it is stated that 76% of American journalists are white, and some 8% of the remaining minority are Hispanic or otherwise Latino (Tomasik, et al., para 14).

Interestingly, according to Figure B2, the age spread appears to be rather evenly distributed. Even though there is some skewing towards younger demographics, with 39% in both the 20-29 and 30-39 range, this would seem to be close to accurate in reality. A 2026 survey conducted by the Reynolds Center for Business Journalism cited that of 325 participants, the average age range was around 40.5 years old for the profession (Culey, para. 8). Of note that isn't displayed in either of these figures is that a decent part of the lower age brackets have images that take place in warzones compared to the later age brackets, whose components mostly take place in cities or protests within those areas.

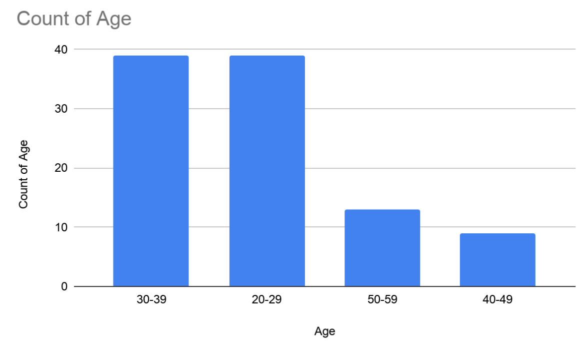

Figure B2: Age Distribution

Discussion

In general, it is fair to presume that the AI, although largely trained off of real datasets, may be being influenced by something else outside of its own judgement. While in reality there are about equal male proportions of journalists compared to female journalists, there may be some influence from popular fictitious media that shapes public perceptions despite reality. Although many major outliers do exist, many major representatives of the profession do wind up being female. Additionally, the outward facing appearance of most female journalists, specifically those on television (who are heavily depicted over traditional paper journalists) are typically seen with conventionally attractive features, whether they be natural, makeup, or otherwise. As a result, the AI may have, in some way, factored this into its generations, and thus created a more largely female generation array, with a hefty amount of them adhering to extreme beauty standards. Aside from that, the overwhelming trend of journalists across most forms, from newspaper to television, being white historically may also be another "subconscious" sect that causes the image generation to display mostly White-European demographics over any other.

Section C: Public Defender

The third part of this study is focused on the Image AI generation of how an average Public Defender is like in the eyes of AI. The findings we discover is that AI seems to always have a bias on every type of job that we research in this study. The study showed that AI was generating more female public defenders rather than male public defenders images. Specifically, 10% of images were male public defenders and 90% of the images were female public defenders. These percentages are from results of 100 images generated by Gemini AI. The pie chart below depicts the data, with the red minority congregation referring to the proportion of male public defenders, and the blue majority representing the proportion of female public defenders in comparison.

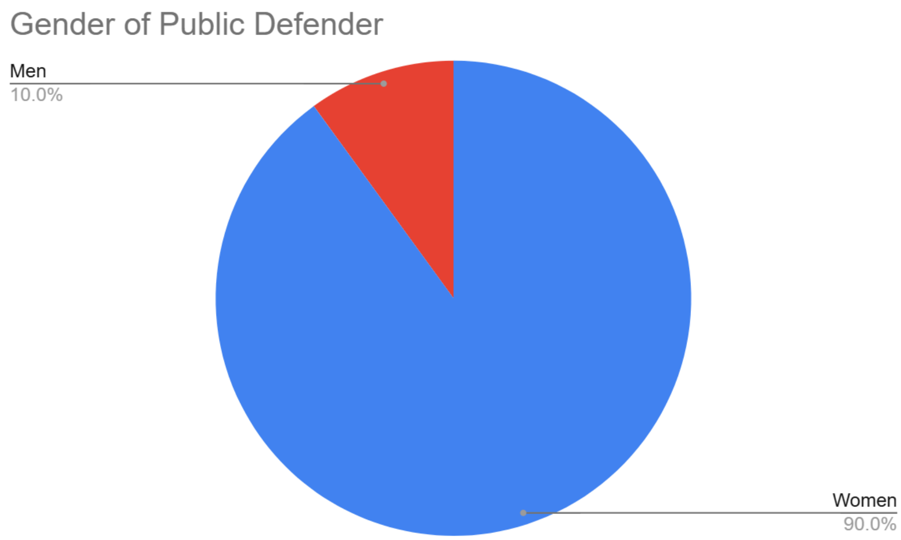

Figure C1: Gender

Based on Figure C2 it shows the variety of ages from the 100 images that we have generated from the AI. This chart is a bar graph that shows 7 different ranges of ages. From the bar graph it shows that 30-35 years old has the highest amount from the images and 45-50 years old are the least amount from the data displayed on the right.

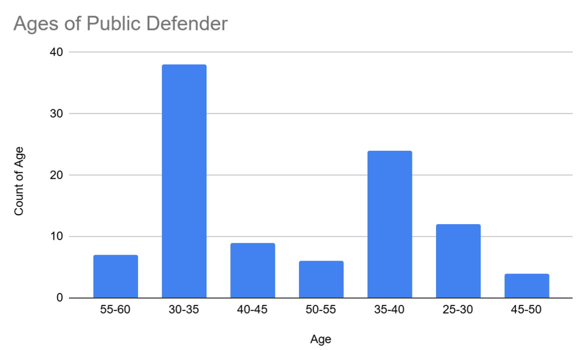

Figure C2: Ages

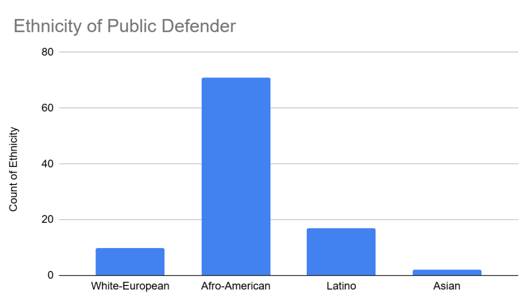

Figure C3: Ethnicities

Figure C3 congregates the data of the ethnicity from the 100 images generated from Gemini AI. Surprisingly, the results showed that African Americans were shown the most and Asians were shown the least. Data is aggregated on the graph to the left.

Discussion:

The results of C1 were pretty surprising to see but the records show that women are the majority that work as public defenders but not as big a gap shown on the pie chart. The data records shown by Zippia states that 54.9% of public defenders are women and 45.1% of public defenders are men (Zippia et al., 2021). The data also shows that in the early days since 2010, men were the majority but after 11 years it started to shift to women being the majority in this job(Zippia et al., 2021). The data we record actually is pretty the same as what the labor records show, being that the women are the majority in both experiments.

The results of C2 were about the ages of the public defenders, the images that were generated by the AI were mostly women that were around 30-35 years old and 35-40 years old. This result is pretty far off from what the labor records states from Zippa, it shows that the average age for public defenders is 46 years old. Looking into it more, it shows that every race is above 40 years olds and no one is younger than 40 (Zippia et al., 2021).

The last graph, C3 is showing data of the ethnicity of the public defenders. The results are actually surprising compared to the Zippia records because the C3 graph shows that African Americans are the majority but Zippia shows that it was mostly white people. Zippia states that 75.9% is white people, 7.7% is Latino, 6.3% is Asians and 5.4% is African Americans(Zippia et al., 2021). It seems like this study is a complete 180 from the labor records. According to an article from China that was experimenting with AI bias as well, states that Black women had the lowest racial accuracy from the AI generated images (Yang, Yiran; 4 Discussion. 2025). From our experiment, the AI didn't seem to have any difficulty generating images of black people at all. The article also states that people of color were mostly generated as white people regardless of gender and racial contexts of the image (Yang, Yiran; 4 Discussion, 2025).

The AI that we used didn't really do any of that with any race when generating the 100 images.

Lastly, some troubles from getting these images was that it took a while to get all 100 images downloaded and ready to upload to the drive. The files were bigger in mega bytes compared to the other two jobs before this one.

Conclusion

Because Nanonbana is a private model, we as users could not know with certainty whether what data the AI model is using or how a model recognizes and operates in cases where real world prejudice exists or existed. Nevertheless we could find the tendencies and patterns that an AI model follows. This is the reason why it causes serious issues beyond the individual inconvenience. For the case of Senator images there is a tendency to follow the archetype of the American politician, which displays the beauty stands for male and women in the United States throughout the 100 images generated. Also the diversity in high strata of the political life of the United States is represented as a minimally diverse group with images that significantly more frequently represent a White-European person. Public perception or otherwise what is most seen of a profession can cause its proponents to only be viewed in one way or another. In the case of journalism, the outward appearance in most televised (and thus more often seen) media often has conventionally attractive female individuals being seen in the position. Additionally, fictional media often depicts similar female journalists. As a possible result, the generations from the AI may have been influenced by that existing media, and thus created more female depictions of journalists compared to males.

By the information that we had gathered in this study, we have discovered that AI displayed a bias in different layers in the three professions. First , the hypothesis that male individuals will be overrepresented in professional roles was wrong, instead generative AI is to display, a majority of time, female individuals in the categories of journalist, senator and public defender. Second, we thought that there would be a diversity between the races for each job but for the majority of them, only one or two races were predominantly featured, White-Europeans especially, with Afro-Americans in second for only the Public Defender position. The information in the images often does not reflect the actual demographics of either age, gender, ethnicity, or all three for any of the given professions. In this regard, the hypothesis has been fulfilled, in that the AI's image generation has been skewed most likely by its own dataset that the images it generates do not mirror reality.

References

Mao, A., Barve, S., Melissa Shi, J., Juneja, P., & Saha, K. (2026). *Social stereotypes in AI text-to-image generation* (Vol. 6). AI and Ethics. https://doi.org/10.1007/s43681-026-01146-8

Brudnick, I. A., & Manning, J. E. (2022, October 13). *Women in Congress: Statistics and brief overview* (CRS Report No. R43244). Congressional Research Service. https://www.congress.gov/crs-product/R43244

United States Senate. (n.d.). *Women senators*. Retrieved July 18, 2026, from https://www.senate.gov/senators/ListofWomenSenators.htm

Zhou, M., Abhishek, V., Derdenger, T., Kim, J., & Srinivasan, K. (2024). Bias in Generative AI. *arXiv*. https://doi.org/10.48550/arxiv.2403.02726 Cited by: 171

Zippia. (2026, July 16). *Public defender jobs near me - 626 jobs.* **https://www.zippia.com/public-defender-jobs/jobs/**

Yang, Y. (2025). Racial bias in AI-generated images. AI & Society, 40(7), 5425–5437. https://doi-org.ccny-proxy1.libr.ccny.cuny.edu/10.1007/s00146-025-02282-1

Simonite, T. (October 24, 2019). *A Health Care Algorithm Offered Less Care to Black Patients.* Wired. https://www.wired.com/story/how-algorithm-favored-whites-over-blacks-health-care/

Jemielity, S. (October 28, 2019). *Health care prediction algorithm biased against black patients, study finds.* UChicago News. https://news.uchicago.edu/story/health-care-prediction-algorithm-biased-against-black-patients-study-finds

Culey, J. (June 24, 2026). *Hiring, raises show strong demand for business journalists: Reynolds Center survey.* The Reynolds Center for Business Journalism https://businessjournalism.org/2026/06/2026-salary-survey/

Tomasik, E, Gottfried, J. (April 4, 2023). *U.S. Journalists' beats vary widely by gender and other factors.* Pew Research Center https://www.pewresearch.org/short-reads/2023/04/04/us-journalists-beats-vary-widely-by-gender-and-other-factors/

Appendix

1. Google Drive

   https://drive.google.com/drive/folders/1CGICZfxny7-1BfZsyxMxXjTLIPeexBhT

2. Image Sheet

   https://docs.google.com/spreadsheets/d/1-5FNX6eeYGLaVP9BFFkxtFyZWjeBB2RwhLKkQa-9Ay0/edit?gid=0#gid=0