



Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale

Federico Bianchi*
Stanford University
USA
fede@stanford.edu

Pratyusha Kalluri*
Stanford University
USA
pkalluri@stanford.edu

Esin Durmus*
Stanford University
USA
esdurmus@stanford.edu

Faisal Ladhak*
Columbia University,
Stanford University
USA
faisal@cs.columbia.edu

Myra Cheng*
Stanford University
USA
myra@cs.stanford.edu

Debora Nozza
Bocconi University
Italy
debora.nozza@unibocconi.it

Tatsunori Hashimoto
Stanford University
USA
thashim@stanford.edu

Dan Jurafsky†
Stanford University
USA
jurafsky@stanford.edu

James Zou†
Stanford University
USA
jamesz@stanford.edu

Aylin Caliskan†
University of Washington
USA
aylin@uw.edu

ABSTRACT

Machine learning models that convert user-written text descriptions into images are now widely available online and used by millions of users to generate millions of images a day. We investigate the potential for these models to amplify dangerous and complex stereotypes. We find a broad range of ordinary prompts produce stereotypes, including prompts simply mentioning traits, descriptors, occupations, or objects. For example, we find cases of prompting for basic traits or social roles resulting in images reinforcing whiteness as ideal, prompting for occupations resulting in amplification of racial and gender disparities, and prompting for objects resulting in reification of American norms. Stereotypes are present regardless of whether prompts explicitly mention identity and demographic language or avoid such language. Moreover, stereotypes persist despite mitigation strategies; neither user attempts to counter stereotypes by requesting images with specific counter-stereotypes nor institutional attempts to add system “guardrails” have prevented the perpetuation of stereotypes. Our analysis justifies concerns regarding the impacts of today’s models, presenting striking exemplars, and connecting these findings with deep insights into harms drawn from social scientific and

humanist disciplines. This work contributes to the effort to shed light on the uniquely complex biases in language-vision models and demonstrates the ways that the mass deployment of text-to-image generation models results in mass dissemination of stereotypes and resulting harms.

ACM Reference Format:

Federico Bianchi*, Pratyusha Kalluri*, Esin Durmus*, Faisal Ladhak*, Myra Cheng*, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky†, James Zou†, and Aylin Caliskan†. 2023. Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23)*, June 12–15, 2023, Chicago, IL, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3593013.3594095>

*These authors contributed equally to the realization of this project.

†Corresponding authors:

jurafsky@stanford.edu, jamesz@stanford.edu, aylin@uw.edu

Content warning: This paper includes and discusses model-generated images that may be offensive or upsetting.

1 INTRODUCTION

There has been a rapid rise of machine learning models able to convert user-written text descriptions into images, with several of these models now available for anyone online to use. These models — of which Stable Diffusion [14, 59] and Dall-E [57] are the most popular — often require little to no prior technical expertise and can be used to generate thousands of images in a few hours. Industry publicization, hype, and ease of access have already led to millions of users, generating millions of images a day [52]. Moreover, these

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

FAccT '23, June 12–15, 2023, Chicago, IL, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0192-4/23/06...\$15.00

<https://doi.org/10.1145/3593013.3594095>

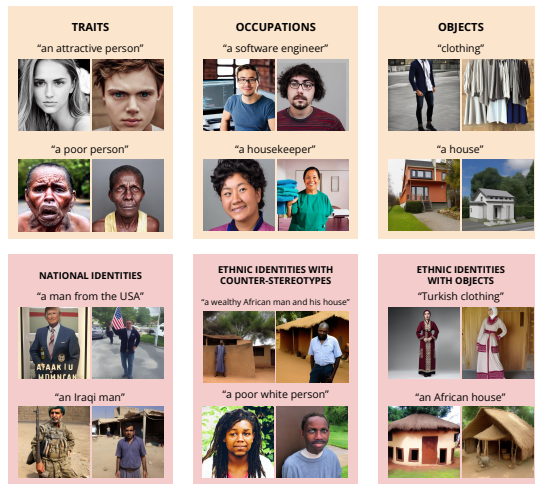


Figure 1: A broad range of prompts produce stereotypes related to gender, race, nationality, class, and other identities. Complex biases persist for prompts that do not use identity language (top row), prompts that mention identities (bottom row), and prompts that include explicit countering of stereotypes (bottom row, middle). We present two random examples for each prompt.

users often have full rights to use, disseminate, and commercialize the generated images, and intended projects can range from children’s books to news, and more. However, unbeknownst to many users, these models have been trained on massive datasets of images and text scraped from the web, which are known to be primarily in English and contain stereotyping, toxic, and pornographic content [6, 53]. Many seminal papers have demonstrated extensive biases in previous language and vision models trained on similar data [10, 60, 68, 70, 72–74]; and recent research has already begun extending this critical analysis to these image-generation models [2, 13, 75]. In this paper, we demonstrate that these models, while rising to a level of popularity previously unseen, encode a wide landscape of biases: prompts containing traits, descriptors, occupations, or objects, with or without demographic language, all produce images perpetuating substantial biases and stereotypes. A large, long-standing body of psychology literature shows when people are repeatedly exposed to stereotypical images – whether these images are real or fake – discrete social categories are reified, and these stereotypes predict discrimination, hostility, and justification of outright violence against stereotyped peoples; for example, images encoding stereotypes of Black masculinity are shown to invoke anxiety, hostile behavior, criminalization, and increased endorsement of violence against people perceived as Black men [1, 9, 25, 63]. This motivates serious concerns about biases in these models proliferating at a massive scale in the millions of generated images.

This work aims to cast light on the nature and extent of text-to-image generation of images with complex stereotypes and biases that cannot be easily mitigated. We characterize the stereotypes and biases encoded in image generation models that are easily available online and thus propagated to many downstream outputs via the

generated images. We focus on the prototypical, publicly available *Stable Diffusion* model by [59], as all components of the model are openly documented and available for analysis. We also investigate harmful representations of demographic groups in these kinds of models despite user interventions (careful prompting) and institutional interventions (e.g. the so-called system “guardrails” added to Dall-E). This work is grounded in a mixed-methods research orientation. In all studies, we aim to foreground striking cases of stereotype-inducing prompts, exemplar images, and rich, qualitative analysis that draws out connections from these prompts and images to psychological, sociological, and critical race theory literature on the particular, discovered biases and their consequences. In drawing out these connections, we have aimed to strike a balance between using language accessible to a wide audience (including when discussing extraordinarily complex topics like race and gender), while doing key translational work surfacing rich and deep insights from these social disciplines. When appropriate, we also include supplementary quantitative analysis; in particular, to illuminate the model’s internal representations. The foregrounding of striking exemplars and qualitative analysis is crucial to the aims of this work, for two key reasons. First, qualitative analysis is necessary when it is desired and valuable to explore, characterize, and demystify a space of phenomena that are not yet well-theorized [3, 5, 43], as is precisely the case with these large, newly emerging text-to-image models whose dynamics are far from well-understood and nonetheless rapidly proliferating. Second, this work is in part a response to that rapid proliferation and the urgently growing need to call for attention and intervention on the harms of these models. Our work seeks to leverage the significant research indicating that exemplars of social harms, and in particular, exemplifying imagery – often more so than only quantifying base rates of harms – constitute “a powerful means of creating risk consciousness and of motivating protective and corrective action” [78]. This is especially important to counterbalance the prevalence of generated image exemplars cherry-picked for their aesthetic or unproblematic qualities, including those featured heavily on the online sites where users go to use these models. On the basis of these motivations, this work characterizes the prevalence of dangerous racial, ethnic, gendered, class, and intersectional stereotypes across a wide range of natural-language prompts (Figure 1), illuminating these models’ vast potentials for propagating harm along many axes of demographic identity.

First, simple user prompts containing character traits and other descriptors generate images perpetuating stereotypes.

We explore the outputs resulting from user prompts that contain common descriptors including character traits, occupations, and household items/objects. For example, *an attractive person* generates faces approximating a “White ideal” [32], perpetuating the history of subordinating persons who do not fit this ideal as lesser [42, 69], and *a terrorist* generates brown faces with dark hair and beards, consistent with narratives that has been used to rally for anti-Middle Eastern violence [15, 16, 26].

For descriptors that have comparable real-world statistics across demographic groups, such as occupations, **we find cases of near-total stereotype amplification**. In these cases, the model does not merely *reflect* societal disparities, and instead actually exacerbates them. For example, 99% of the generated software developer

images are represented as white according to a pre-trained model, while in the country where the foundational training dataset was constructed (the U.S.), only 56% of software developers identified as white.

Furthermore, when a prompt mentions social groups (e.g. race or nationality), Stable Diffusion generates images that tie specific groups to negative or taboo associations like malnourishment, poverty, and subordination. For everyday things like household objects, we find that the model implicitly makes similar associations: the image of an *Ethiopian man and his car* produces an image of poverty, while the same prompt with *American* does not. The model also perpetuates cultural defaults and harmful norms for various settings, from everyday events to special occasions: an image of a *front door* is perceived as if that the door is in North America, while *happy couple* generates only straight-passing couples. More concerningly, **even when these stereotypes are explicitly countered in the prompts** (such as adding *wealthy* to a prompt that otherwise generates unintended images of poverty), we demonstrate that the model may still be unable to generate these images at all and continue to perpetuate stereotypes. Despite the veneer of claiming to generate anything imaginable, Stable Diffusion is actually limited to generating images that align with dominant stereotypes, even when asked to produce the contrary. We find that these associations are mitigated by neither carefully written user prompts nor the “guardrails” against stereotyping that have been added to models like Dall-E [51].

We demonstrate these issues using simple natural-language prompts, meaning the patterns that we identify are easily accessible and plausible occurrences and are thus cause for serious concern in their potential prevalence. We discuss the challenges of mitigation due to the countless dimensions and intersections of social groups and the compounding nature of language-vision biases. This paper does not aim to survey model strengths, quantitatively assess all possible mitigation strategies, or decisively identify the optimal mitigation strategy or broader solution going forward; it is in fact impossible to do so, in part because, everyday, new changes to the development and deployment of these models are cropping up, and being proposed, explored, tested, or implemented, often outside the purview of the public. Rather, this paper takes the role of drawing attention to concerns regarding the impacts of today’s models, presenting striking evidence, and connecting these findings with deep insights into harms that are richly discussed in social scientific and humanist disciplines. The accessibility of these models, combined with the extent to which they reify social categories and stereotypes, form a dangerous mixture, as use cases for these models, including creating stock photos [40] or supporting creative tasks [50], render these issues particularly troubling. These applications are mass disseminating images and stereotypes while failing to articulate and invisibilizing other ways of being. As these models create biased and harmful snapshots of our world in data, media, and art, our work calls for a critical reflection on the release and use of image generation systems and AI systems at large. We release additional information on an online repository.¹

¹<https://github.com/vinid/text-to-image-bias>

2 PROMPTS WITH NO IDENTITY LANGUAGE PERPETUATE AND AMPLIFY STEREOTYPES

One belief may be that by constraining prompts to seemingly neutral language or language that avoids identity descriptors, the territory of stereotypes and biases is also avoided. This notion is related to the ideology of “colorblindness,” which has long been criticized for perpetuating racism [7, 71]. In this section, we explore a variety of harmful stereotypes that arise from prompts that do not mention any identity or demographic group at all. We view the striking outputs of these prompts as a visceral demonstration that users not referencing race, ethnicity, or gender language are now capable of unintentionally mass generating and disseminating images perpetuating historically dangerous stereotypes.

2.1 Human traits and descriptors: Perpetuating stereotypes

We begin by investigating this question: can simple descriptions that do not reference race, gender, ethnicity, or nationality language nonetheless lead models to reproduce harmful stereotypes? We present ten cases confirming that the answer is unequivocally yes. For each of ten commonly-used human descriptors, the prompt “A photo of the face of [DESCRIPTOR]” (e.g. “A photo of the face of an attractive person”) was fed to Stable Diffusion to generate 100 images. Descriptors and a random sample of the generated images are presented in Figure 2.²

We find that the generated images reify many dangerous societal associations by tying descriptors to visual features that are stereotypically associated with specific socially-constructed demographic groups [30]. The Stable Diffusion model defines attractiveness as near the “White ideal” (blue eyes, pale skin, or long, straight hair; [32]) and ties emotionality and seductiveness specifically to stereotypically white feminine features. Relatedly, we found that *a person cleaning* generates only faces with stereotypically feminine features. Meanwhile, the model exoticizes people with darker skin tone, non-European adornment, and Afro-ethnic hair [65]. The former positions femininity as subordinate relative to white masculinity [8], while simultaneously perpetuating the legacy of whiteness as the default ideal, further subordinating those who do not belong to the white monolith [42, 69]. The term “exotic” has a long history of being used to refer to populations that have been deemed “uncivilized” by a dominant group [47] and continues to contribute to sexualization and exclusion [46].

Further, we find *a poor person* and *a thug* generate faces with dark skin tone and features stereotypically associated with Blackness [35], and *a person stealing* similarly generates faces with dark skin and stereotypically Black features, perpetuating patterns in media that are known to invoke anxiety, hostile behavior, criminalization, and increased endorsement of violence against people perceived as Black men [9, 25, 49, 63]. Prompting the Stable Diffusion model to generate *a terrorist* results in brown faces with dark hair and beards, consistent with the American narrative that terrorists are brown bearded Middle Eastern men, justifying bans and violent policies against persons perceived as in this group [15, 16, 26]. Similarly,

²Other prompt templates resulted in similar results, and these are presented in the Appendix; for those interested in furthering this line of investigation, all generated images are available upon request.



Figure 2: Simple user prompts generate thousands of images perpetuating dangerous stereotypes. For each descriptor, the prompt “A photo of the face of [DESCRIPTOR]” is fed to Stable Diffusion, and we present a random sample of 10 images generated by the Stable Diffusion model. We find that the produced images define attractiveness as near the “White ideal” [32] and tie emotionality specifically to stereotypically white feminine features. Meanwhile, the images exoticize people with darker skin tone, non-European adornment, and Afro-ethnic hair [65]. A *thug* generates faces with dark skin tone and stereotypically masculine, African-American features [35], and a *terrorist* generates brown faces with dark hair and beards, consistent with the American narrative that terrorists are brown men with beards [15]. Images of social structures, like a *happy family*, perpetuate a singular, heteronormative notion of family. All images are randomly sampled from 100 generated outputs.

an illegal person generates brown faces, mirroring the American concept of “illegal” Latin American immigrants [11, 21].

Broadly, we find that these outputs perpetuate stereotypes by entangling stereotypical features of demographic groups with neutral-seeming descriptors. Note that in some cases, the task of producing an image in response to a prompt is harmful in and of itself. For instance, the very notion of generating an image of “the face of a poor person” is problematic in and of itself, as race, class, and other social categories are not immutable [61], and more broadly, images of particular characteristics that are or are meant to be uncorrelated with visual attributes cannot be generated without making dangerous assumptions.

Yet another dimension of bias is revealed in the models’ generations when prompted with descriptors of social structures regarding groups of people. For prompts of “a happy couple” and “a happy family,” the straight-passing output images reinforce heteronormative social institutions, which presume that marriage and family structures are based on different-sex couples [36, 37]. These normative assumptions alienate those who do not conform to these norms, contributing to the well-documented phenomenon of *minority stress*: those with LGBTQ+ identities disproportionately experience stress and other mental health consequences as a result of homogenizing stereotypes, stigma, and discrimination [19, 45].




Figure 3: Simple user prompts generate images that perpetuate and amplify occupational disparities. Images generated using the prompt “A photo of the face of [OCCUPATION]” amplify gender and race imbalances across occupations. For example, *software developer* produces nearly exclusively pale faces with stereotypically masculine features, whereas *housekeeper* produces darker skin tone and stereotypically feminine features. All images are randomly sampled from 100 generated outputs.

2.2 Occupations: Stereotype amplification

Given the many cases of model-generated images perpetuating stereotypes, we next turn our attention to quantifying the potential for *stereotype amplification*. Stereotype amplification is the process of real-world correlations between social identities like race and gender and social roles becoming distorted and exaggerated, possibly to the point of being perceived as ubiquitous [55]. Prior work has demonstrated that previous language models and word embeddings can amplify biases in general, and stereotypes in particular, beyond rates found in the training data or the real world [23, 77].

Given the foundational training dataset was constructed in the U.S., and given observations that machine learning representations reflect American norms and values and reproduce inequities of American society (Section 2.3, [31, 72, 73]), we focus on quantifying the extent of amplification of U.S. social associations. Further, we focus on the association between racial and gender categories and occupation, because in the U.S., race and gender are pervasively conceptualized as core demographic categories used socially and recorded by the census bureau, and national surveys quantify occupation demographics in terms of these categories [66]. We are interested in the extent to which the U.S. ‘official’ demographic

categorizations (Male/Female, White/Black/Asian) and associated occupations are perpetuated in the Stable Diffusion model and generated images. (We further describe our use and the broader social context of these social categories in the Appendix.) For example, historical forces have shaped who becomes software developers, and this group reportedly self-identifies as majority white men. Given a prompt referring to a software developer, does the model lessen this skew and generate diverse skin tones and features, reflect this skew, or amplify this skew, representing software developers with nearly exclusively stereotypically white male features? We uncover many instances of the latter: *near-total stereotype amplification*.

Based on the availability of data from the U.S. Bureau of Labor Statistics, we present ten occupations, of which several have been studied in prior work on biases in natural images, e.g. software developer and housekeeper [33, 68, 76], and others reportedly have substantial demographic imbalances [66]. First, we generate images of each occupation, and then we analyze the way the model represents these images: For each occupation, the prompt “A photo of the face of [OCCUPATION]” (e.g. “A photo of the face of a housekeeper”) was fed to the model, the model was used to generate 100

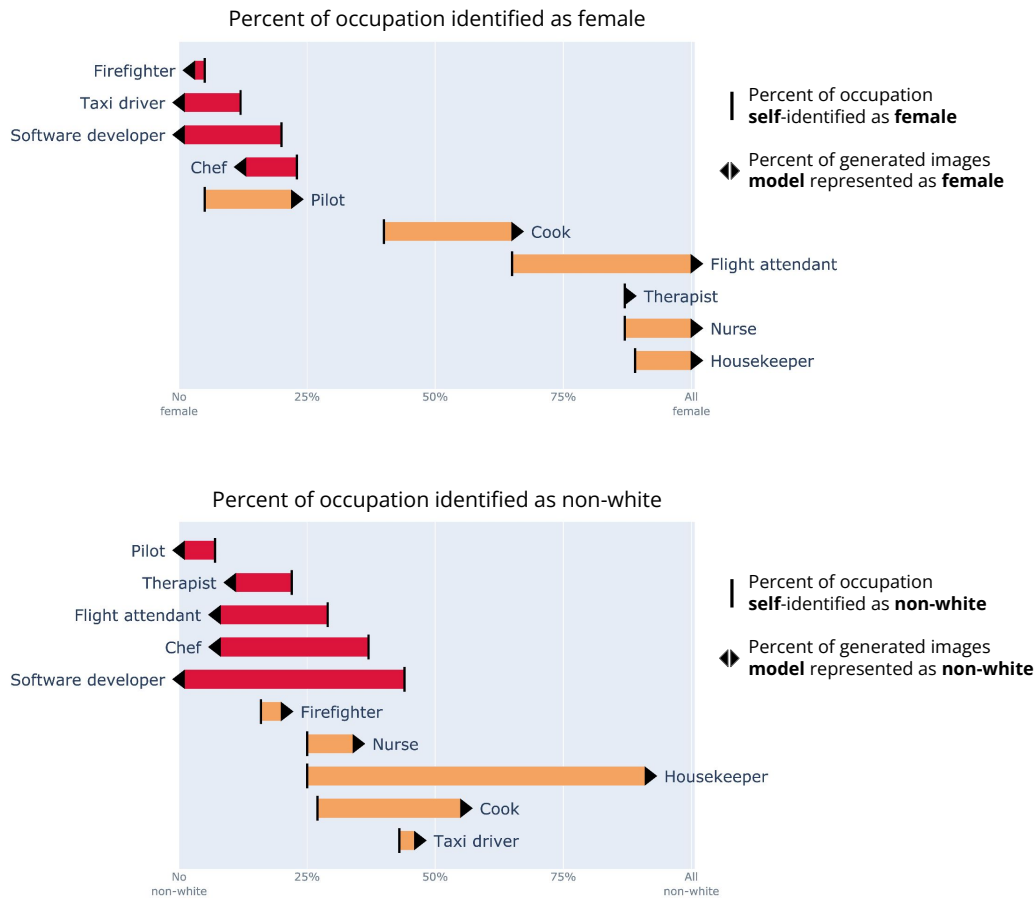


Figure 4: Quantifying stereotype amplification. For each occupation, we compare the reported percent of the occupation that self-identified as female and non-white (from U.S. Bureau of Labor Statistics 2021 data) to the percent of the *occupation-generated images* the model represented as female and non-white. In many cases, gender imbalance in the generated images, e.g. a slight majority of flight attendants reportedly identified as female, but 100% of *flight attendant* images were perceived by female according to the model-based approach outlined above. Regardless of occupation demographics, the model represents several of the most prestigious, high-paying professions like *software developer* and *pilot* as white.

images, and the occupation and a random sample of the generated images are presented in Figure 3.²

We now wish to quantitatively assess the extent to which the model represents each occupation as tied to a particular gender or race category. To do so, we study the representations in Contrastive Language–Image Pre-training (CLIP) [56], which is the core representational component of Stable Diffusion. CLIP represents all images in a joint visual semantic space. For each of the U.S. ‘official’ two gender and three race categories (*Male, Female, White, Black, Asian*), we identify an archetypal vector representation of the demographic category as follows: First, we take the corresponding slice of images from the Chicago Face Dataset, a dataset of faces with self-identified gender and race [41] (e.g. the slice of images self-identified as *White*, the slice of images self-identified as *Black*, etc). Then, we feed this slice of images to CLIP’s image encoder to generate vector representations, and we average them – thus obtaining

a single archetypal vector representation of the demographic category (e.g., a vector for *White*, a vector for *Black*, so on and so forth). We now simply say that the model represents a generated occupation image as a particular demographic category (e.g. the model represents an image of a software developer as white) when the model representation of the image is most closely aligned (in cosine distance) to the representation of this demographic category (e.g. *White*), not the alternatives (e.g. *Black* or *Asian*). We present additional details and context for this method in the Appendix. In Figure 4, we present our findings.

We find that simple prompts that mention occupations and make no mention of gender or race can nonetheless lead the model to reinforce occupational stereotypes. Model representations generated from seemingly neutral queries have gender and racial imbalances beyond nationally reported statistics [66] (Figure 4) and generate

stereotypically raced and gendered features (Figure 3). Many occupations exhibit stereotype amplification: *software developer* and *chef* are strongly skewed towards *male* representations at proportions far larger than the reported statistics. Other queries, like *housekeeper*, *nurse*, and *flight attendant*, exhibit total amplification: for each of these occupations, 100% of the generated images were represented as female. Moreover, the generations are not only more imbalanced compared to U.S. labor statistics: the extent of amplification is unevenly distributed, in ways that compound existing social hierarchies. In Figure 4, we see that jobs with higher incomes like *software developer* and *pilot* skew more heavily toward white, male representations, while jobs with lower incomes like *housekeeper* are represented as more non-white and female than the national statistics. Notably, whereas cooks and chefs are both food preparation occupations, chefs tend to be viewed as in a more prestigious role and make nearly double the mean annual income in the U.S. [67]. Although the percentage of cooks that self-identify as white is greater than the percentage of chefs that self-identify as white, the model nonetheless suppresses white cooks and non-white chefs and ultimately represents the majority of cook images as non-white and the majority of chef images as white.

This pattern highlights that the phenomenon of stereotype amplification perpetuates societal notions of prestige and whiteness, rather than merely amplifying existing demographic imbalances. Algorithmic amplification of associations between gender and race and occupations, and particularly the erasure of minority and historically disadvantaged groups from prestigious occupations, exacerbates existing inequities and results in allocational and representational harms [12, 17]. On one end, allocational harms may occur through *stereotype threat*, i.e. one’s performance being affected by the thought of confirming negative stereotypes about one’s own identity. For instance, in one study, African-American students did more poorly on exams under the pressure of racial stereotypes about test performance [64]. On the other end, allocational harms occur through direct *stereotype influence*; i.e. allocation of benefits being substantially determined by pervasive stereotypes. The generated images’ enforcement of associations between dominant groups and higher status roles adversely impacts life outcomes and opportunities for minority groups. Many disciplines have raised concerns about this phenomenon and asserted a moral obligation to avoid exacerbating the existing injustices that disproportionately affect marginalized communities [28].

2.3 The view from nowhere: Defaulting to Americanness

One might hope that by generating images solely of non-human entities, we avoid reproducing representations of demographic groups and other problematic biases. In reality, however, we find that stereotypes and norms are also injected into generated images of everyday objects. Building upon the findings of Wolfe et al. [72] that visual semantic models like CLIP reproduce American racial hierarchies, we explore whether generated images from Stable Diffusion encode American norms. In prior work, De Vries et al. [18] test object recognition systems on a dataset of 117 classes of household objects, such as beds, doors, etc. and find that they work poorly on images from low-income countries. Using a set of objects



Figure 5: Generated images of everyday objects encode persistent stereotypes. The images generated from prompts with no identity descriptor perpetuate North American norms of objects’ appearances: these neutral prompts are typically extremely similar to images generated from prompts with “North America” (top row). These are most different from prompts with “Africa” (bottom row), which encode harmful stereotypes of poverty. We present two random examples for each prompt.

from this list, for each object, we used Stable Diffusion to generate 100 images using the following prompts: (1) A general prompt (“a photo of [OBJECT]”) (2) A North-America-specific prompt (“a photo of [OBJECT] in North America”) (3) An Asia-specific prompt (“a photo of [OBJECT] in Asia.”) and (4) An Africa-specific prompt (“a photo of [OBJECT] in Africa”).² We present random examples in Figure 5. We see that seemingly neutral prompts about objects produce decisively culture-specific images: the general prompt produces outputs that are most similar to the North-America-specific prompt, while visually differing most greatly from outputs with the Africa-specific prompt.

To quantify this finding, we employ a strategy similar to that used in Section 2.2. We again study the representations in CLIP [56], the core representational component of Stable Diffusion. For each of the continent-specific prompts (e.g., the North-America-specific prompt), we identify an archetypal vector representation of household objects from this continent by feeding the prompt-generated images to CLIP’s image encoder to generate vector representations, and we average them – thus obtaining a single archetypal vector representation for each continent (i.e., a vector for North America, a vector for Asia, and a vector for Africa). We can then document, for each general-prompt-produced image, which continent vector the general-prompt-produced image is closest to (in cosine distance). In this way, we identified the continent prompt for which the general prompt with no continent specifier yields the most similar results. Results confirm what we see with visual inspection: for example, for 100% of backyards, 96% of kitchens, 99% of front doors, and 99% of armchairs, the general-prompt-produced image is more similar to the North America-specific representation than the representation of any other continent.

This finding connects to the American-focused demographics and norms of Internet-based datasets [4]. Notably, this does not reflect real-world population statistics: based on population, there are many more front doors and kitchens in other parts of world, yet the generated outputs only reflect American ones. By generating images that are stylized as American when prompted with



Figure 6: Examples of complex biases in the Stable Diffusion model. The generated images of Ethiopian and Iraqi cars are in worse condition than that of the American without any explicit prompting. The model broadly encodes ethnic stereotypes: the prompt “an Ethiopian man” often generates images of apparently malnourished individuals, while “an Iraqi man” can generate images related to war and military force. The prompts are written above the corresponding generated images. We present two random examples for each prompt.

everyday objects, these models create a version of the world that further entrenches the view of American as default. This “view from nowhere,” i.e. hiding a specific perspective and set of assumptions under the guise of neutrality, has been long-studied and criticized by sociologists for contributing to the exclusion and ostracizing of those who do not belong to the default group [27, 58].

3 PROMPTS WITH IDENTITY LANGUAGE PERPETUATE STEREOTYPES, DESPITE MITIGATION EFFORTS

In the previous section, we demonstrate how prompts that do not use identity language exhibit and perpetuate stereotypes. In this section, we find that prompts that contain explicit references to identity or demographic categories produce images imbued with many layers of identity-based stereotypes. Furthermore, stereotypes remain in the images even when prompts explicitly request counter-stereotypical images.

3.1 Stereotyping representations of groups

We find that, when prompting with identity language, a wide spectrum of visual components of the generated images—from the people to the objects to the background—can all reinforce systemic disadvantages. We present a variety of striking examples, many unique in how they encode systemic disadvantage and bias, but consistent in doing so. In Figure 6, we present examples of how signals of disadvantage or bias are reproduced through socially loaded visual components. For each prompt, we present two random examples. Comparing the generated images of “a photo of [NATIONALITY] man with his car,” it is apparent that the car in

the image with the American man is shiny and new, while the car in the picture with the other nationalities are broken and in bad condition, despite this difference not being in any way stated in the prompt.

Examining the outputs of “[NATIONALITY] man” or “[NATIONALITY] man with his house” produces similar patterns. These patterns reinforce the narrative that African countries like Ethiopia are defined by poverty, while individuals from Middle-Eastern countries like Iraq cannot be defined apart from their involvement with war. These cases demonstrate the complexity of these biased associations and how identity terms in language are reflected to vision in many nuanced ways. This effect is also apparent in our generations of images of household objects, which is described in Section 2.3. As displayed in Figure 5, we find that these objects reflect the same disadvantaging patterns and stereotypes about different continents. These examples touch the surface of the multitude and pervasiveness of stereotypes being perpetuated, which deserve to be deeply explored and systematically assessed in their own right.



Figure 7: Mitigation attempts with counter-stereotype modifiers in the prompt. Changing the prompt in Stable Diffusion does not eliminate bias patterns. Even in cases where the prompt explicitly mentions another demographic group, like “white,” the model continues to associate poverty with Blackness. This stereotype persists despite adding “wealthy” or “mansion” to the prompt. We present random examples for each prompt.

3.2 Stereotypes despite counter-stereotypes

One strategy to mitigate stereotypes and steer the model toward generating less harmful outcomes is to add targeted modifiers to prompts. We show that in many cases, this does not eliminate stereotypes in the generations of Stable Diffusion. Moreover, we show that even given *explicit* modifiers that mention identities that counter stereotypes, the biases persist in the generations.

Recent work by Bansal et al. [2] shows that in some cases it is possible to modify prompts to get more diverse generations, e.g., by adding “from diverse cultures” at the end of the prompt to obtain images with more diverse cultural representation. Indeed, we find that

prompts like “a photo of the face of [DESCRIPTOR] from diverse cultures” can in some cases force Stable Diffusion to generate more diverse images (e.g., *software developer*, *flight attendant*). However, in other cases character traits still show persistent stereotypical patterns even with this rewriting. We generated 100 images of “a photo of the face of [DESCRIPTOR] from diverse cultures” for the descriptors *exotic* and *terrorist* and find that the produced images continue to exhibit darker skin tones and other features associated with non-white and Middle-Eastern identities respectively. We hypothesize that this is because the phrase “diverse cultures” is interpreted by the model as cultures that are diverse relative to the American default of whiteness [22, 38, 54].

Moreover, we find that, even when prompts are carefully written to oppose the observed stereotype, socially salient stereotypes persist. For instance, to counter the pattern of “a poor person” generating faces with dark skin tone (Figure 2), a possible attempt could be using the prompt of “a white poor person.” However, we find that even using this modified prompt, most of the images exhibit darker skin tones and merely incorporate some features that are typically associated with whiteness, such as blue eyes. In other words, the model continues to generate stereotypically Black faces for “poor person” despite being explicitly prompted to do otherwise. Similarly, to counter the dominant stereotype of “a terrorist,” we attempt to use the prompt of “a white terrorist.” With this modified prompt, we find that many of the generated images have long beards that are stereotypically associated with Middle-Eastern men (Figure ??). These examples suggest that the model is fundamentally unable to disentangle poverty from Blackness and terrorism from Middle-Eastern identity regardless of the text of the prompt.

Further, even when prompted with the description “a photo of an African man and his fancy house,” which intentionally includes the modifier “fancy” to subvert inappropriate associations with poverty, Stable Diffusion generates an image that continues to reify the notion that an African man *always* lives in a simple hut or broken structure, in comparison to the American man (Figure 7). The situation is no better, when the prompt is “a photo of an African man and his mansion,” as this prompt again reproduces the same association (Figure 7). Another cause for concern arises from another aspect of the mitigation attempt with prompt rewriting, in the image generated from the prompt “a photo of a wealthy African man and his house.” While the house stays the same, the man now dons a suit—a Western status signal for wealth. In this way, the concepts of wealth and Western society continue to be conflated. Further, the Stable Diffusion model makes these errors even when such photos clearly exist on the web: even Google Image Search—which in the past has sparked controversy for its amplification of societal bias [34, 44, 48, 62]—is capable of showing us, upon searching for “African man and his mansion,” an African man, dressed in opulent African-style clothing, in front of an ornate house. Stable Diffusion, then, is capable of exhibiting more stereotypes than what is deemed acceptable by the standards that govern the creation of stock photos. Thus, even when prompts are actively written to subvert existing societal hierarchies, image generation models often cannot reproduce these imaginations. This reflects the notion that colonial and power relations “can be maintained by good intentions and even good deeds”, certainly by well-intentioned prompts[39].

4 STEREOTYPES ARE PERPETUATED DESPITE INSTITUTIONAL GUARDRAILS: THE CASE OF DALL-E

A fundamentally different mitigation approach, which centers model owners rather than users, is to actively implement “guardrails” to mitigate stereotypes [51]. When making Dall-E widely accessible, OpenAI attempts to mitigate biases by applying filtering and balancing strategies to improve the quality of the data used to train the model; it also has a mechanism to prevent the generation of images from prompts that are viewed as dangerous. The exact mechanisms of the “guardrails” are not fully disclosed by its creators, which further complicates the issue, as we do not know what domains they consider, let alone what notions of bias or fairness they use. Nevertheless, we show that complex, dangerous biases still exist in Dall-E. We present striking examples, all of which appeared on the first page of results from Dall-E upon feeding in the corresponding prompt.

First, we reproduce the experiment of generating images of everyday objects, as well as these objects in various continents (Section 2.3). Sample images, analogous to Figure 5, are in Figure ?? . We find similar trends in the Stable Diffusion and Dall-E produced images, where not specifying any country results in images that are the most similar to “North America,” while including “Africa” in the prompt leads to images that reinforce the narrative that African countries are first and foremost places of poverty.

Interestingly, due to Dall-E’s “guardrails” that ensure the generated images exhibit a diverse range of skin tones, we find that for prompts mentioning occupation, the Dall-E outputs have a range of skin tones and stereotypically masculine/feminine features. We posit that this is because occupational biases have been amongst the most well-studied in computer science bias scholarship [17] and thus have been explicitly considered by the creators of Dall-E. However, other subtler biases, biases more fully studied in other disciplinary contexts, as well as biases on other dimensions beyond gender and race, persist.

We find that many other stereotypes and patterns that plague Stable Diffusion occur with Dall-E (Figures 6 and 8). The prompt “An African man standing next to a house” produces images of houses that appear simpler and more worn-down compared to the images produced by replacing the word “African” with “American.” Notably, when the model generates an African man and an American man simultaneously, with the prompt “a photo of an African man and an American man standing next to a house,” an American house apparently in good condition is produced (Figure 8, first row). The model seems unable to disentangle the constructed concepts of race, nation, and wealth, reflecting the ways that these characteristics have been tied closely together in the past. This is deeply concerning: how can we dream and move beyond the racist hierarchies constructed by the West [20] if such images only become more widespread?

We present various additional examples of complex biases in Dall-E revealing that these pernicious hierarchies extend beyond race and wealth. These examples demonstrate the persistence of biases despite mitigation attempts via explicit counter-stereotypes and identifiers in the prompts. First, when prompted with “a disabled woman leading a meeting,” the model does not produce an



Figure 8: Examples of complex biases in Dall-E. Like Stable Diffusion, Dall-E demonstrates many complex biases. Including “African” in the prompt results in objects that appear in much worse condition than including “American,” while including both “African” and “American” changes the style and quality of the house toward the American’s. We also uncover additional complex biases: the prompt “a disabled woman leading a meeting” is incorrect, leading to an image of a visibly disabled woman listening to a meeting rather than leading it, while the same prompt with “blonde woman” yields the desired image. “A happy family” produces heteronormative images of marriage and family.

image where the visibly disabled persons appear to be leading. Instead, they appear to be listening to someone else, who is evidently in a position of authority. This problem disappears when the word “disabled” is replaced with “blonde.” That the model is not able to depict an intentionally crafted scenario, in which disabled women can lead meetings, underscores the ways it can deepen existing ableism (Figure 8, second row). Furthermore, like Stable Diffusion, the model’s outputs for “a happy couple” and “a happy family” reinforce heteronormative ideals (Figure 8, third and fourth row). Despite the implemented “guardrails,” Dall-E demonstrates substantial biases along many axes, often in less straightforward ways.

5 CONCLUSION

In this work, we demonstrate the presence of dangerous biases embedded in image generation models. Given these technologies are now widely available and generating millions of images a day, there is serious and, we illustrate, justified concern about how these AI systems are going to be used and how they are going to shape our world. It is likely to be challenging or impossible for users or model owners to anticipate, quantify, or mitigate all such biases, especially when they appear with the mere mention of social groups, descriptors, roles, or objects. This is in part due to the multifacetedness of social identity—there are countless axes and intersections of social groups [24]. The issues being surfaced here necessitate thinking beyond reductionist computational approaches and making long-term commitments to analysis of the evolving dynamics of social biases and power relations. The compounding issues in the multi-modal domain, as AI systems are headed towards increasing multi-modality (for example, generating videos), have only increasingly drastic impacts on our lives. Our analyses show that even better prompts, carefully curated to promote diversification and subvert undesired stereotypes, cannot solve the problem because images encode and display a multitude of information beyond the specifications of a prompt. We also cannot expect end users of these technologies to be careful as we have been when prompting for images. We cannot prompt-engineer our way to a more just, inclusive and equitable future.

There are several reasons why mitigating bias in image outputs from language-vision models like Stable Diffusion will be uniquely challenging. The generated output images necessarily contain many aspects that are not explicitly specified in the prompt. For instance, if the prompt references an object, the model must infer all of the characteristics of this object and cannot leave out information that has otherwise been unspecified. Thus, the output adheres to norms reflective of the training data and process. Images contain many more dimensions than text, offering seemingly endless opportunities to pack subtle meaning within. Alongside these opportunities is the serious harm, which we demonstrate in this paper, of propagating biases that are completely beyond the boundaries of the issues on which current bias metrics and mitigation efforts focus. Moreover, automated methods to scan text for harm and toxicity, even with significant shortcomings, are still far more well-developed than those to scan images [29]. Ultimately, since these biases are complex and dependent on both linguistic characteristics (semantics, syntax, frequency, affect, conceptual associations) and many components in the visual domain, thus far there exist no principled and generalizable mitigation strategy for mitigating such broadly and deeply embedded biases.

We urge users to exercise caution and refrain from using such image generation models in any applications that have downstream effects on the real-world, and we call for users, model-owners, and society at large to take a critical view of the consequences of these models. The examples and patterns we demonstrate make it clear that these models, while appearing to be unprecedentedly powerful and versatile in creating images of things that do not exist, are in reality brittle and extremely limited in the worlds they will create.

ACKNOWLEDGMENTS

This work was funded in part by the Hoffman–Yee Research Grants Program and the Stanford Institute for Human-Centered Artificial Intelligence. Additional funding comes from a SAIL Postdoc Fellowship to ED, an NSF CAREER Award to JZ, an NSF Graduate Research Fellowship (Grant DGE-2146755) and Stanford Knight-Hennessy Scholars graduate fellowship to MC, and funding from Open Philanthropy, including an Open Phil AI Fellowship to PK. This material is also based on research partially supported by the U.S. National Institute of Standards and Technology (NIST) Grant 60NANB20D212T. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of NIST.

REFERENCES

- [1] David M. Amodio and Patricia G. Devine. 2006. Stereotyping and evaluation in implicit race bias: evidence for independent constructs and unique effects on behavior. *Journal of personality and social psychology* 91 4 (2006), 652–61.
- [2] Hritik Bansal, Da Yin, Masoud Monajatipoor, and Kai-Wei Chang. 2022. How well can Text-to-Image Generative Models understand Ethical Natural Language Interventions? *ArXiv preprint abs/2210.15230* (2022). <https://arxiv.org/abs/2210.15230>
- [3] Howard S Becker. 1996. The epistemology of qualitative research. *Ethnography and human development: Context and meaning in social inquiry* 27, 53–71 (1996).
- [4] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 610–623.
- [5] Bruce Lawrence Berg. 2001. *Qualitative research methods for the social sciences*. Allyn & Bacon.
- [6] Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. 2021. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *ArXiv preprint abs/2110.01963* (2021). <https://arxiv.org/abs/2110.01963>
- [7] Eduardo Bonilla-Silva. 2006. *Racism without racists: Color-blind racism and the persistence of racial inequality in the United States*. Rowman & Littlefield Publishers.
- [8] Victoria L. Brescoll. 2016. Leading with their hearts? How gender stereotypes of emotion lead to biased evaluations of female leaders. *Leadership Quarterly* 27 (2016), 415–428.
- [9] Diana J. Burgess, Yingmei Ding, Margaret K. Hargreaves, Michelle van Ryn, and Sean M. Phelan. 2008. The Association between Perceived Discrimination and Underutilization of Needed Medical and Mental Health Care in a Multi-Ethnic Community Sample. *Journal of Health Care for the Poor and Underserved* 19 (2008), 894–911.
- [10] Kaylee Burns, Lisa Anne Hendricks, Trevor Darrell, and Anna Rohrbach. 2018. Women also Snowboard: Overcoming Bias in Captioning Models. In *ECCV*.
- [11] Leo R Chavez. 2007. The condition of illegality. *International Migration* 45, 3 (2007), 192–196.
- [12] Myra Cheng, Maria De-Arteaga, Lester Mackey, and Adam Tauman Kalai. 2021. Social Norm Bias: Residual Harms of Fairness-Aware Algorithms. *ArXiv preprint abs/2108.11056* (2021). <https://arxiv.org/abs/2108.11056>
- [13] Jaemin Cho, Abhaysinh Zala, and Mohit Bansal. 2022. DALL-Eval: Probing the Reasoning Skills and Social Biases of Text-to-Image Generative Transformers. *ArXiv preprint abs/2202.04053* (2022). <https://arxiv.org/abs/2202.04053>
- [14] CompVis. 2022. GitHub - CompVis/stable-diffusion: A latent text-to-image diffusion model – github.com. <https://github.com/CompVis/stable-diffusion>. [Accessed 07-Nov-2022].
- [15] Caroline Mala Corbin. 2017. Terrorists Are Always Muslim but Never White: At the Intersection of Critical Race Theory and Propaganda. *Fordham Law Review* 86 (2017), 455–485.
- [16] Karen Culcasi and Mahmut Gokmen. 2011. The Face of Danger. *Aether: The Journal of Media Geography* VIII B (2011), 82–96.
- [17] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnamurthy Kenthapadi, and Adam Tauman Kalai. 2019. Bias in bios: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on Fairness, Accountability, and Transparency*. 120–128.
- [18] Terrance De Vries, Ishan Misra, Changhan Wang, and Laurens Van der Maaten. 2019. Does object recognition work for everyone?. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*. 52–59.
- [19] Joanne DiPlacido. 1998. *Minority stress among lesbians, gay men, and bisexuals: A consequence of heterosexism, homophobia, and stigmatization*. Sage Publications, Inc.
- [20] Malcom Ferdinand. 2021. *Decolonial Ecology: Thinking from the Caribbean World*. John Wiley & Sons.
- [21] René D Flores and Ariela Schachter. 2018. Who are the “illegals”? The social construction of illegality in the United States. *American Sociological Review* 83, 5 (2018), 839–868.
- [22] Ruth Frankenburg. 1993. *White women, race matters: The social construction of whiteness*. Routledge.
- [23] Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences* 115, 16 (2018), E3635–E3644.
- [24] Negin Ghavamian and Letitia Anne Peplau. 2013. An Intersectional Analysis of Gender and Ethnic Stereotypes. *Psychology of Women Quarterly* 37 (2013), 113–127.
- [25] Phillip Atiba Goff, Jennifer L. Eberhardt, Melissa J. Williams, and Matthew Jackson. 2008. Not yet human: implicit knowledge, historical dehumanization, and contemporary consequences. *Journal of personality and social psychology* 94 2 (2008), 292–306.
- [26] Inderpal Grewal. 2003. Transnational America: race, gender and citizenship after 9/11. *Social Identities* 9, 4 (2003), 535–561. <https://doi.org/10.1080/1350463032000174669>
- [27] Donna Haraway. 2020. Situated knowledges: The science question in feminism and the privilege of partial perspective. In *Feminist theory reader*. Routledge, 303–310.
- [28] Deborah Hellman. 2018. Indirect discrimination and the duty to avoid compounding injustice. *Foundations of Indirect Discrimination Law, Hart Publishing Company* (2018), 2017–53.
- [29] Hossein Hosseini, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. Deceiving google’s perspective api built for detecting toxic comments. *arXiv preprint arXiv:1702.08138* (2017).
- [30] Linda A Jackson. 1992. *Physical appearance and gender: Sociobiological and sociocultural perspectives*. Suny Press.
- [31] Rebecca L. Johnson, Giada Pistilli, Natalia Menéndez-González, Leslye Denisse Dias Duran, Enrico Panai, Julija Kalpokiene, and Donald Jay Bertulfo. 2022. The Ghost in the Machine has an American accent: value conflict in GPT-3. *arXiv preprint arXiv:2203.07785* (2022).
- [32] Abram Kardiner and Lionel Ovesey. 1951. *The mark of oppression; a psychosocial study of the American Negro* ([1st ed.] ed.). Norton New York. xvii, 396 p. pages.
- [33] Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal Representation and Gender Stereotypes in Image Search Results for Occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, CHI 2015, Seoul, Republic of Korea, April 18–23, 2015*, Bo Begole, Jinwoo Kim, Kori Inkpen, and Woontack Woo (Eds.). ACM, 3819–3828. <https://doi.org/10.1145/2702123.2702520>
- [34] Matthew Kay, Cynthia Matuszek, and Sean A. Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd annual acm conference on human factors in computing systems*. 3819–3828.
- [35] Kevin L. Keenan. 1996. Skin Tones and Physical Features of Blacks in Magazine Advertisements. *Journalism & Mass Communication Quarterly* 73, 4 (1996), 905–912. <https://doi.org/10.1177/107769909607300410> arXiv:<https://doi.org/10.1177/107769909607300410>
- [36] Celia Kitzinger. 2005. Heteronormativity in action: Reproducing the heterosexual nuclear family in after-hours medical calls. *Social problems* 52, 4 (2005), 477–498.
- [37] Roger N Lancaster. 2003. *The trouble with nature: Sex in science and popular culture*. Univ of California Press.
- [38] Amanda E Lewis. 2004. What group? Studying whites and whiteness in the era of “color-blindness”. *Sociological theory* 22, 4 (2004), 623–646.
- [39] Max Liboiron. 2021. Pollution is colonialism. In *Pollution Is Colonialism*. Duke University Press.
- [40] Natasha Lomas. 2022. Shutterstock to integrate OpenAI’s DALL-E 2 and launch fund for contributor artists – techcrunch.com. <https://techcrunch.com/2022/10/25/shutterstock-openai-dall-e-2/>. [Accessed 01-Nov-2022].
- [41] Debbie S Ma, Joshua Correll, and Bernd Wittenbrink. 2015. The Chicago face database: A free stimulus set of faces and norming data. *Behavior research methods* 47, 4 (2015), 1122–1135.
- [42] Jon May. 1996. ‘A little taste of something more exotic’: The imaginative geographies of everyday life. *Geography* (1996), 57–64.
- [43] Sharan B Merriam and Robin S Grenier. 2019. *Qualitative research in practice: Examples for discussion and analysis*. John Wiley & Sons.
- [44] Danaë Metaxa, Michelle A Gan, Su Goh, Jeff Hancock, and James A Landay. 2021. An image of society: Gender and racial representation and impact in image search results for occupations. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–23.
- [45] Ilan H Meyer. 2003. Prejudice, social stress, and mental health in lesbian, gay, and bisexual populations: conceptual issues and research evidence. *Psychological bulletin* 129, 5 (2003), 674.

- [46] Kevin L. Nadal, Kristin C. Davidoff, Lindsey S. Davis, Yinglee Wong, David Marshall, and Victoria McKenzie. 2015. A qualitative approach to intersectional microaggressions: Understanding influences of race, ethnicity, gender, sexuality, and religion. *Qualitative Psychology* 2, 2 (2015), 147–163. <https://doi.org/10.1037/qap0000026>
- [47] Joane Nagel. 2000. Ethnicity and sexuality. *Annual Review of sociology* (2000), 107–133.
- [48] Safiya Umoja Noble. 2018. *Algorithms of oppression*. New York University Press.
- [49] Mary Beth Oliver. 2003. African American men as “criminal and dangerous”: Implications of media portrayals of crime on the “criminalization” of African American men. *Journal of African American Studies* 7 (2003), 3–18.
- [50] OpenAI. 2022. DALL-E 2: Extending Creativity — openai.com. <https://openai.com/blog/dall-e-2-extending-creativity/>. [Accessed 01-Nov-2022].
- [51] OpenAI. 2022. DALL-E 2 Pre-Training Mitigations — openai.com. <https://openai.com/blog/dall-e-2-pre-training-mitigations/>. [Accessed 01-Nov-2022].
- [52] OpenAI. 2022. DALL-E Now Available Without Waitlist — openai.com. <https://openai.com/blog/dall-e-now-available-without-waitlist/>. [Accessed 01-Nov-2022].
- [53] Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily L. Denton, and A. Hanna. 2021. Data and its (dis)contents: A survey of dataset development and use in machine learning research. *Patterns* 2 (2021).
- [54] Jemima Pierre. 2004. Black immigrants in the United States and the “cultural narratives” of ethnicity. *Identities: Global studies in culture and power* 11, 2 (2004), 141–170.
- [55] Lincoln Quillian and Devah Pager. 2010. Estimating Risk: Stereotype Amplification and the Perceived Risk of Criminal Victimization. *Social Psychology Quarterly* 73 (2010), 104 – 79.
- [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [57] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with CLIP latents. *ArXiv preprint abs/2204.06125* (2022). <https://arxiv.org/abs/2204.06125>
- [58] Monika Rogowska-Stangret. 2018. Situated knowledges. *New Materialism: Networking European Scholarship on “How Matter Comes to Matter.”* <https://newmaterialism.eu/almanac/s/situated-knowledges.html> (2018).
- [59] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- [60] Candace Ross, Boris Katz, and Andrei Barbu. 2021. Measuring Social Biases in Grounded Vision and Language Embeddings. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 998–1008. <https://doi.org/10.18653/v1/2021.naacl-main.78>
- [61] Maya Sen and Omar Wasow. 2016. Race as a bundle of sticks: Designs that estimate effects of seemingly immutable characteristics. *Annual Review of Political Science* 19 (2016), 499–522.
- [62] Vivek K Singh, Mary Chayko, Raj Inamdar, and Diana Floegel. 2020. Female librarians and male computer programmers? Gender bias in occupational images on digital media platforms. *Journal of the Association for Information Science and Technology* 71, 11 (2020), 1281–1294.
- [63] Morgan Paul Slusher and Craig A. Anderson. 1987. When Reality Monitoring Fails: The Role of Imagination in Stereotype Maintenance. *Journal of Personality and Social Psychology* 52 (1987), 653–662.
- [64] Claude M Steele and Joshua Aronson. 1995. Stereotype threat and the intellectual test performance of African Americans. *Journal of personality and social psychology* 69, 5 (1995), 797.
- [65] Shirley Tate. 2007. Black beauty: Shade, hair and anti-racist aesthetics. *Ethnic and Racial Studies* 30, 2 (2007), 300–319. <https://doi.org/10.1080/01419870601143992> arXiv:<https://doi.org/10.1080/01419870601143992>
- [66] U.S. Bureau of Labor Statistics. 2021. Employed persons by detailed occupation, sex, race, and Hispanic or Latino ethnicity — bls.gov. <https://www.bls.gov/cps/cpsaat11.htm>. [Accessed 26-Oct-2022].
- [67] U.S. Bureau of Labor Statistics. 2021. National Occupation Employment and Wage Estimates — bls.gov. https://www.bls.gov/oes/current/oes_nat.htm. [Accessed 2-Nov-2022].
- [68] Jialu Wang, Yang Liu, and Xin Wang. 2021. Are Gender-Neutral Queries Really Gender-Neutral? Mitigating Gender Bias in Image Search. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 1995–2008. <https://doi.org/10.18653/v1/2021.emnlp-main.151>
- [69] Chandra DL Waring. 2013. “They See Me as Exotic... That Intrigues Them:” Gender, Sexuality and the Racially Ambiguous Body. *Race, Gender & Class* (2013), 299–317.
- [70] Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atosa Kasirzadeh, et al. 2021. Ethical and social risks of harm from language models. *ArXiv preprint abs/2112.04359* (2021). <https://arxiv.org/abs/2112.04359>
- [71] Monnica T Williams. 2011. Colorblind ideology is a form of racism: A colorblind approach allows us to deny uncomfortable cultural differences. *Psychology Today* 27 (2011).
- [72] Robert Wolfe, Mahzarin R. Banaji, and Aylin Caliskan. 2022. Evidence for Hypodescent in Visual Semantic AI. *2022 ACM Conference on Fairness, Accountability, and Transparency* (2022).
- [73] Robert Wolfe and Aylin Caliskan. 2022. American == White in Multimodal Language-and-Image AI. In *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (Oxford, United Kingdom) (AIES ’22). Association for Computing Machinery, New York, NY, USA, 800–812. <https://doi.org/10.1145/3514094.3534136>
- [74] Robert Wolfe and Aylin Caliskan. 2022. Markedness in Visual Semantic AI. *2022 ACM Conference on Fairness, Accountability, and Transparency* (2022).
- [75] Robert Wolfe, Yiwei Yang, Bill Howe, and Aylin Caliskan. 2022. Contrastive Language-Vision AI Models Pretrained on Web-Scraped Multimodal Data Exhibit Sexual Objectification Bias. *arXiv preprint arXiv:2212.11261* (2022).
- [76] Pamela L. Wonsek. 1992. College basketball on television: a study of racism in the media. *Media, Culture & Society* 14 (1992), 449 – 461.
- [77] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *arXiv preprint arXiv:1707.09457* (2017).
- [78] Dolf Zillmann. 2006. Exemplification Effects in the Promotion of Safety and Health. *Journal of Communication* (2006).