

Interpretations, Representations, and Stereotypes of Caste within Text-to-Image Generators

Sourojit Ghosh

University of Washington, Seattle
ghosh100@uw.edu

Abstract

The surge in the popularity of text-to-image generators (T2Is) has been matched by extensive research into ensuring fairness and equitable outcomes, with a focus on how they impact society. However, such work has typically focused on globally-experienced identities or centered Western contexts. In this paper, we address interpretations, representations, and stereotypes surrounding a tragically underexplored context in T2I research: *caste*. We examine how the T2I Stable Diffusion displays people of various castes, and what professions they are depicted as performing. Generating 100 images per prompt, we perform CLIP-cosine similarity comparisons with default depictions of an ‘Indian person’ by Stable Diffusion, and explore patterns of similarity. Our findings reveal how Stable Diffusion outputs perpetuate systems of ‘castelessness’, equating Indianness with high-castes and depicting caste-oppressed identities with markers of poverty. In particular, we note the stereotyping and representational harm towards the historically-marginalized Dalits, prominently depicted as living in rural areas and always at protests. Our findings underscore a need for a caste-aware approach towards T2I design, and we conclude with design recommendations.

1 Introduction

With the swift rise in prominence of text-to-image generators (T2Is) since 2021 across public and commercial usage with a valuation of almost US \$45 billion this year (Lawlor and Chang 2023), it is increasingly important to study the societal and ethical impacts of the outputs of such tools. As sociotechnical systems where societal conditions and hierarchies have a direct impact on their outputs (Shelby et al. 2022), T2Is have the immense power of shaping viewers’ opinions and sense of reality by propagating depictions – problematic or otherwise – of groups or individuals, at scale (Hall 1997; Qadri et al. 2023). This power typically produces disproportionately unfair outputs for groups traditionally marginalized in society, as demonstrated by researchers focusing on a wide range of harms such as allocational or representational harms (Barocas et al. 2017) along the lines of aspects of identity including gender (e.g., Ghosh and Caliskan 2023b), race (e.g. Bianchi et al. 2023), and disability (e.g. Mack et al. 2024). As crucial as such research

is, these works have often centered around perspectives and stereotypes centered around the Global North, with little focus on identities from non-Western parts of the world.

In this paper, we focus on one such identity: **caste**. As an identity and system of discrimination prevalent in South Asia and applicable to almost half of the global population, caste remains woefully unexplored in the context of representations and harms embedded within T2I outputs. To start bridging this gap, we explore the research question:

RQ: *How is caste interpreted and represented within T2I outputs and what (if any) stereotypes are embedded within outputs of prompts which mention caste?*

We focus on the T2I *Stable Diffusion*, whose high popularity and open-source nature has made it the go-to model to study how harms are propagated at scale (e.g., Ghosh and Caliskan 2023b; Luccioni et al. 2023). We focus on a series of prompts around caste-based identity and stereotypes around occupations, since caste is an identity historically associated with occupation (Zwick-Maitreyi et al. 2018). By comparing sets of 100 Stable Diffusion outputs per prompt using CLIP-cosine similarity – a metric used to compare images returning a score in the 0-1 range where higher score implies higher similarity, previously used for a similar work by Ghosh and Caliskan (2023b) – supplemented by manual qualitative verification, we make two contributions:

- * We demonstrate patterns of ‘casteless’ representations of Indianness (Vaghela, Jackson, and Sengers 2022) within Stable Diffusion outputs. Performing CLIP-cosine similarity comparisons of 100 Stable Diffusion outputs for ‘Indian person’ with those of ‘Indian high-caste person’, as well as outputs for ‘Indian Brahmin person’, ‘Indian Kshatriya person’ and ‘Indian Vaishya person’, we observe how similarity scores fall in the 0.77-0.71 range. Manual comparison of images also substantiates this high similarity, as images seem to portray similar faces. Analogous prompts comparing outputs for ‘Indian person, at work’ with those of ‘Indian high-caste person, at work’, as well as outputs for ‘Indian Brahmin person, at work’, ‘Indian Kshatriya person, at work’ and ‘Indian Vaishya person, at work’, produces CLIP-cosine similarity scores in the 0.76-0.67 range, and manual examination of images also shows a common pattern of people across images mostly being seated at desks/on chairs and

working with either a laptop/keyboard or writing/reading from notebooks. These findings indicate how Stable Diffusion outputs equate the default Indianness with being high-caste, reifying existing power dynamics within India and the diaspora, that relieve caste-privileged individuals from acknowledging their privilege and actively supporting caste-oppressed identities.

* We also show that the casteless representation of Indianness within Stable Diffusion outputs dissipates when portraying historically caste-oppressed identities, as stereotypical depictions emerge. CLIP-cosine similarity comparisons of 100 Stable Diffusion outputs for ‘Indian person’ with those of ‘Indian low-caste person’, as well as outputs for ‘Indian Shudra person’, ‘Indian Dalit person’ and ‘Indian Adivasi person’ produces scores in the 0.63-0.37 range. Manual qualitative verification of these outputs also reveals that low-caste and Shudra individuals are portrayed as shirtlessly standing in farmland or in front of mud houses, while representations of Dalit people often show them in large groups appearing to be on some sort of protest or march. Furthermore, comparing outputs for ‘Indian person, at work’ with those of ‘Indian low-caste person, at work’, as well as outputs for ‘Indian Shudra person, at work’, ‘Indian Dalit person, at work’ and ‘Indian Adivasi person, at work’ yields CLIP-cosine similarity scores in the 0.53-0.44 range. Upon manual verification, we observe that while the ‘Indian person, at work’ is commonly shown seated at desks and typing/writing at work, patterns across the aforementioned prompts show depicted individuals sitting/standing and appearing to be performing work such as bricklaying or gathering wood. These results point to a systematic Othering of and propagating harmful stereotypes towards low-caste identities within Stable Diffusion outputs, targeted strongly towards caste-oppressed and Dalit people.

Through these findings, we demonstrate a pattern within Stable Diffusion outputs that reifies existing power dynamics within communities both in India and the diaspora by producing a “rendering in which upper-caste individuals are able to frame themselves as largely casteless (and meritorious), while lower-caste individuals are seen as still marked by caste.” (Vaghela, Jackson, and Sengers 2022). We conclude with design recommendations for T2Is such as Stable Diffusion to develop improve interpretations of caste and work towards more equitable representation within outputs.

2 Background and Related Work

2.1 T2Is, and Stable Diffusion

Text-to-image generators (T2Is) are generative AI tools that take in text-based prompts from users and provide one or more images as outputs. They are built on top of multimodal large language models (LLMs) such as OpenAI’s GPT series with GPT-4 and GPT-3, and are typically trained upon large datasets of text and images sourced from the Internet.

In this paper, we specifically focus on *Stable Diffusion*: a T2I built upon the vision-language model CLIP to process text prompts and machine translate resultant text embeddings to images (OpenAI 2021; Radford et al. 2021). CLIP

uses picture-caption pairs to learn joint text image embeddings which enhances the semantic space it can represent (Wolfe et al. 2023), compared to representations based on image features alone. Stable Diffusion takes in a text-based prompt, tokenizes it and produces word embeddings through CLIP, and passes it into a UNet noise predictor alongside a random latent noise image to determine how much noise should be reduced from the currently-noisy image. After subtracting this noise, the process is repeated a certain number of times, which can be specified but is set to 50 by default on the web interface¹. After repeating this process, a variational auto-encoder (VAE) neural network is used to convert the latent noise-reduced image into the pixel space, to generate an output. The model is trained on the LAION 5B dataset (Schuhmann et al. 2022), which consists of 5 billion text-image pairs sourced from the open Internet.

Stable Diffusion was launched in 2022 by Stability AI 2022 as “a latent text-to-image diffusion model capable of generating photo-realistic images given any text input, [which] empowers billions of people to create stunning art within seconds”. Its open-source nature and wide popularity to by millions of users globally in personal and commercial contexts has made it a common T2I to study (e.g., Ghosh and Caliskan 2023b; Luccioni et al. 2023), and we too chose this as the T2I in which to examine harmful outputs.

2.2 Harms Caused by T2Is

The evaluation of and research into T2Is has shown that their outputs can cause *harm*: adverse experiences directly or indirectly caused by T2I outputs (Shelby et al. 2022). In the T2I context, the types of harm most commonly considered are *allocative* and *representational harms*. Proposed by Barocas et al. (2017), allocative harms are those where opportunities or resources are withheld from individuals or groups by virtue of their identities, whereas representational harms are those surrounding unfairly-constructed depictions of individuals/groups which may lead viewers to form negative opinions or stereotypes. Representational harms are further sub-categorized by Dev et al. (2020) into five types – *stereotyping*, or the overrepresentation of a set of beliefs about an identity, *disparagement*, or the idea that some groups of people are lesser than others, *dehumanization*, or the practice of treating certain groups of people as less than human, *erasure*, or the lack of representation of groups of people, and *quality of service*, or when a model provides inequitable outcomes for different groups of people. Research into harms caused by the outputs of T2Is (e.g. Ghosh et al. 2024; Gautam, Venkit, and Ghosh 2024; Qadri et al. 2023) has mostly centered around representational harms.

Such research has all, albeit through different approaches and focuses, come towards the general conclusion that the documented harms caused by the outputs of T2Is are usually towards or felt by individuals or groups with one or more identities which have historically been marginalized. This claim centers around the understanding that at their core, ML-driven systems like T2Is rely upon and themselves perform classifications, as they assign a variety of labels

¹<https://stablediffusionweb.com/app/image-generator>

upon individuals and groups. Such systems of classification, Bowker and Star (1999) contend, “give advantage or they give suffering, ... [and] how these choices are made, and how we may think about that invisible matching process, is at the core of the ethical project of this work.” In the context of T2Is, these classifications are often partially or entirely performed by humans within the design pipeline, who do so based on their own subjective perspectives. These subjective perspectives are rooted in individual positionalities, which are informed by their lived experiences within societies that contain historical hierarchies of power and position based on various aspects of identity. Morgan (2018) places these on ‘axes of privilege, domination, and oppression,’ a wheel of line with a common center where each line represent an axis of identity (e.g. gender, race, age) and the end-points of the lines represent the most and least privileged identities along that axis. These axes are intentionally intersecting, recognizing that no individual hosts a single identity and that their lived experiences are shaped by a complex intersection of their various identities (Crenshaw 2017), with individuals who hold multiple identities at the lower end of multiple hierarchies experiencing marginalizations more than those with fewer such low-privilege identities (Collins 1990).

However, research around harms and marginalizations perpetuated by AI typically has had a focus on Global North perspectives e.g., recruiting interviewees primarily from Global North countries, performing comparisons to statistics and stereotypes from such countries like highlighting the replication of specific stereotypes around people of color present in the US (e.g., Benjamin 2019) or studying bias based on US statistics (e.g., Caliskan, Bryson, and Narayanan 2017; Ghosh and Caliskan 2023a). Although such research is undoubtedly important, often missing from conversations around harms caused by the outputs of T2Is are aspects of identity which prominently occur in countries and societies outside of the Global North (Qadri et al. 2023), creating a critical gap in the field. In this paper, we start addressing this gap by focusing on one such aspect: **caste**.

2.3 Caste, and Caste-Based Discrimination

Though the word ‘caste’ originates from the Iberian ‘casta’ meaning ‘lineage/hierarchy’ brought by Portuguese colonizers to India in the 1700s (Bharath 2023), the caste system is known to predate the word by several millennia. Caste, or its native Indian equivalents ‘*jati*’ or ‘*varna*’, is believed to have first emerged in the ancient Hinduism text *Rigveda* around 1500 BC. In it, it is said that humans all originate directly from the body of Lord Brahma the Creator. From His mouth emerge the **Brahmins**, from His arms the **Kshatriyas**, from His thighs the **Vaishyas**, and from His feet the **Shudras**. Beyond these are the **Dalits**, formerly referred to as the ‘untouchables’, and the **Adivasis** or the Indigenous people of India. A visual representation of this hierarchy is shown in Figure 1. The first three castes – Brahmins, Kshatriyas, and Vaishyas – are considered high/upper-caste and collectively called *Savarnas*, whereas Shudras, Dalits and Adivasis are caste-oppressed (Zwick-Maitreyi et al. 2018) *Avarnas*. In more recent history, the latter group is collectively also known as *Bahujans*, popularized by champion of Dalit

rights such Dr. B.R. Ambedkar and Jyotirao Phule. Caste is hereditarily determined at birth, is unchangeable over an individual’s lifetime and affects every aspect of their lives, especially through the practice of caste-based endogamy.

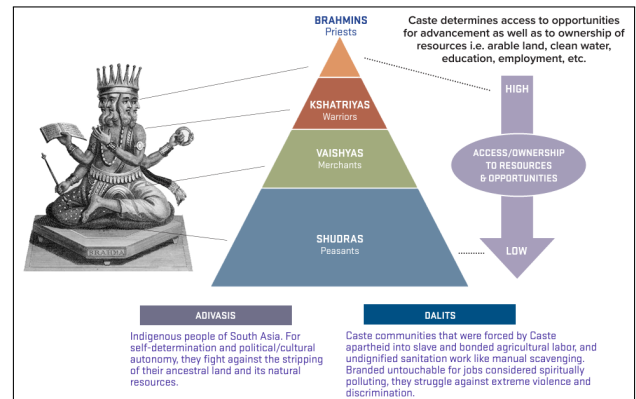


Figure 1: Visualization of the Caste pyramid and socio-religious hierarchy, sourced from Equality Labs (2018).

The association of castes with professions is formalized in a later text, the *Manusmriti* (1-3 AD), forming the bedrock of what we today understand as the ‘caste system’. As shown in Figure 1, the modern caste system is encoded as a hierarchy based on occupations. The rationale was that Brahmin priests are the closest to God and therefore must have the highest power, followed by the Kshatriya rulers who were responsible for law and order, followed by the Vaishya traders who kept the economy going, and then the Shudra workers who formed the backbone of the labor force (Saha 2020). At the very bottom, low enough to not even be assigned a place on the pyramids, are Adivasis who were considered primitive people only fit for farming, and Dalits who performed labor such as sewage and sanitation work. The codification of a profession-based hierarchy doomed Dalits to lifetimes of oppression, relegating them to a sub-human status by labeling them ‘untouchable’ and instituting a belief that they made impure everything they touched.

Such a system of caste-based oppression, the impact of which is felt particularly by Dalits and Shudras, is often referred to as **caste apartheid** (Shetty 1979). The term draws upon the system of apartheid witnessed in South Africa and largely, contends that caste apartheid is comparable to racist practices all over the world, a similarity acknowledged by Dr. Ambedkar in a letter to African-American Civil Rights activist W.E.B DuBois (Kapoor 2003) and the formation of the Dalit Panthers Party inspired by the Black Panther party in the US (Shetty 1979). Some scholars do warn against using ‘racism’ and ‘casteism’ interchangeably, like Wilkerson (2020)’s contention of ‘racism’ to only be about 500 years old starting from the transatlantic trade and very much inspired by the millenia-old caste system. She labels modern-day racism an ‘American invention’ purposely fluid to assign race-based privileges (or lack thereof) to immigrants from different countries, whereas caste is far more rigidly determined at birth as defined in the ancient past.

Indian history is filled with concerted struggle and activism by caste-oppressed groups advocating for equity. Some examples include movements in the 1980s by the Dalit Panthers to reduce casteist harassment in the workplace, nationwide movements in 2016 following the suicide of Dalit scholar Rohith Verma after continual casteist oppression by his University, the 2018 Delhi march when the Indian Supreme Court hearing a case on reducing legal protections for caste-oppressed groups, and numerous Adivasi-led protests in the 1960s and 70s against landowners. Such movements led to the establishment of some protections for caste-oppressed groups. As the Father of the Indian Constitution, Dr. Ambedkar ensured that caste was enshrined as a protected class in the Constitution, and the larger protected categories of **Scheduled Castes (SCs)**, **Scheduled Tribes (STs)** and **Socially and Educationally Backward Classes** who are also known as **Other Backward Classes (OBCs)**. In future years, the government passed legislation such as the 1955 Protection of Civil Rights Act criminalizing the practice of untouchability and the 1989 Prevention of Atrocities Act to provide stronger protections and assign special courts for prosecuting injustices against SCs and STs and measures for providing relief to affected individuals. Furthermore, in recognition of the historical and systematic lack of access of resources for SCs/STs/OBCs, the Indian government employs a system of *reservations*: where the Central/State governments can stipulate that certain percentages of jobs, funding opportunities/ scholarships, or seats in higher educational institutions must be allocated to such communities. Towards affirmative actions for caste-oppressed groups, important to highlight is the work of the 1979 Socially and Educationally Backward Classes Commission (SEBC), also known as the Mandal Commission after the chairperson B.P. Mandal. The Mandal Commission report ensured a collective 49.5% reservation in government and public sector resources for SCs, STs and OBCs, which has undoubtedly been beneficial towards equitable opportunities.

Despite these protections, caste apartheid remains prevalent, operating in both overt and invisibilized ways. In July 2024, students at India's Jawaharlal Nehru University found the casteist slogan (translated) 'Dalit, leave India' written on their hostel walls (Raj 2024b). Around the same time, the Karnataka state government proposed the cancellation the equivalent of \$1.6bn previously allocated to support SC/ST students pursuing PhDs in foreign universities, but later reversed the decision after large-scale student protests erupted (R 2024). Investigative media outlet *The Wire* also brought to the Supreme Court in 2024 their findings from an extensive project on prisoner conditions in various Indian states, uncovering a series of caste-discriminatory practices which the Supreme Court termed 'most disturbing' (Shantha 2024). These are some recent examples of caste apartheid in India, amidst a sea of daily practices at the macro and micro levels.

It is a common but incorrect conception that the caste system and the effects of caste apartheid are confined to India and should not be considered a global issue. Firstly, though India is one of the largest countries with the longest histories of caste apartheid, caste-based systems of discrimination also exist elsewhere across Asia and the rest of the world:

Balinese caste systems in Indonesia divide people into Brahmins/priests, Satrias/ knights, Wesias/businesspeople, and Sudras/workers (Boon 1977); the 'hukou' system in China is considered to similar to caste (Ho 2003); Nigeria practiced a caste system where the 'Osu' person was ostracized based on heredity until it was recently abolished in 2018 (Abia, Amalu, and Ariche 2021), to name a few examples. Systems of social stratification linked to occupation and assigned at birth, where certain groups are considered almost sub-human, are thus common across the world. Secondly, even within the Indian context, the caste system firmly extends to the diaspora. Most Indian immigrants to countries such as the US and the UK are Savarna which lead to "religious and political institutions being created mostly by "upper" caste immigrants who established "upper" caste Hindu culture as the norm." (Zwick-Maitreyi et al. 2018). This severely Others low-caste immigrants and places them in a doubly precarious position as being an outsider through both their immigration status and caste. Surveying the Indian diaspora within the US, Equality Labs (2018) reported that Dalits live in a constant fear of being 'outed' to their coworkers/ friends which would result in discrimination and for people whose castes are known to others, a majority experienced physical/verbal abuse or workplace discrimination.

Despite this, caste remains an unprotected category in global laws. In countries such as the US and UK where significant portions of populations are Indian immigrants or citizens of Indian descent, affirmative action laws do not offer protection against caste-based discrimination. In the US, the city of Seattle is the first and only one to add caste to its anti-discrimination laws (Rajvanshi 2023), a landmark win referred to by leading Dalit rights activist Thenmozhi Soundararajan as "a win centuries in the making," (Kaur 2023) achieved in 2023 after many years of efforts since a 2015 march by the All India Dalit Women's Rights Forum. Though some universities like California State University and those within the University of California system also ban caste discrimination, these policies exist in a very small number of contexts and workplaces which unfortunately leave open the conditions for rampant casteist discrimination within the Indian diaspora. In other contexts, the UK's Equality Act of 2010 only allows Ministers to discretionarily order caste to be treated as an aspect of race and does not explicitly ban caste-based discrimination (Harad 2023). The United Nations' legal framework² contains several instruments to protect against discrimination around identities such as race/ethnicity, sexual orientation, gender, disability status, immigration, Indigeneity, and religion, but not caste.

Representations of caste in the outputs of machine learning tools has been rarely explored, with the only examples being Qadri et al. (2023)'s brief focus on disempowerment of Dalit communities in T2Is, Tiwari et al. (2022)'s study on language models associating negative sentiments with Dalits in Hindi and Tamil, and Dammu et al. (2024)'s findings on how LLMs perpetuate extreme negative opinions on caste. Our work intends to bridge this critical gap in the field.

²<https://www.un.org/ruleoflaw/thematic-areas/human-rights/equality-and-non-discrimination/>

3 Methods

3.1 Prompt Formation and Image Generation

To explore our research question, we adopt a strategy similar to Ghosh and Caliskan (2023b) and use the prompt ‘a front-facing image of a person from India’ to use as a baseline against which to compare other results. We chose the construction ‘... a person from India’ as opposed to ‘...an Indian person’ because the latter can be confounded with ‘American Indian’ referring to populations Indigenous to the US.

We use caste-embedded prompts of ‘a front-facing image of a high-caste person from India’ and ‘a front-facing image of a low-caste person from India,’ using the hyphenated version to keep the qualifier ‘low’/‘high’ associated with ‘caste.’ We also created 6 prompts around the castes shown in Figure 1: ‘a front-facing image of a _____ person from India’, filling in blanks with the 6 castes ‘Brahmin’, ‘Kshatriya’, ‘Vaishya’, ‘Shudra’, ‘Dalit’, and ‘Adivasi’ e.g., ‘a front-facing image of a Dalit person from India’, etc. We collectively refer to these prompts as **Caste-Only prompts**.

Given the association of caste with occupation (described in Section 2.3), we created occupation-based prompts starting with a baseline (‘a front-facing image of a person from India, at work’) and then prompts embedding high- or low-caste, as well as 5 prompts with the labels Brahmin, Kshatriya, Vaishya, Shudra, Dalit, and Adivasi e.g., ‘a front-facing image of a high-caste person from India, at work’, ‘a front-facing image of a Dalit person from India, at work’. Commas are added into prompts to not imply that a person’s caste identity only exists while they are at work. We collectively refer to these prompts as **Caste-Occupation prompts**.

Though Ghosh and Caliskan (2023b) generated 50 images per prompt in their study, more recent research (Du et al. 2024) has seen that number rise to 100. We thus decided to generate 100 images per prompt, hoping that it would show larger variation in results and that patterns documented over 100 images per prompt would be stronger than those over 50. Images were generated on the most updated open-source version of Stable Diffusion (v2.1) available at the time of this writing. We use a self-developed codebase for generating a large volume of images and, to both mirror user experience on the Stable Diffusion interface which does not ask users to specify seeds and in keeping with prior work (Ghosh and Caliskan 2023b), do not assign a deterministic seed for image generation. Data was collected in January 2024.

Similar to the approach of Ghosh and Caliskan (2023b), we too refer to prompts in this paper in a shorted format: the prompt ‘a front-facing image of a person from India’ is shortened to ‘Indian person’, ‘a front-facing image of a Dalit person from India’ is shortened to ‘Indian Dalit person’, ‘a front-facing image of a caste-privileged person from India’ is shortened to ‘Indian caste-privileged person’, ‘a front-facing image of a person from India, at work’ is shortened to ‘Indian person at work’, ‘a front-facing image of a Dalit person from India, at work’ is shortened to ‘Indian Dalit person at work’, and ‘a front-facing image of a caste-privileged person from India, at work’ is shortened to ‘Indian caste-privileged person at work’, etc.

3.2 Analysis Techniques

Our primary method of analyzing Stable Diffusion outputs for the aforementioned prompts is through CLIP-cosine similarity. This approach has been used for image comparison and, most closely related to our work, specifically to evaluate Stable Diffusion outputs (Ghosh and Caliskan 2023b).

Cosine similarity is a technique of comparing two vectors, represented by a score between 0 and 1 where 0 indicates total dissimilarity and 1 implies the vectors are the same (Singhal et al. 2001). In our approach, we use CLIP to obtain embeddings for each image within the output set of 100 images per prompt, and vectorize those embeddings to compute cosine similarity comparisons akin to Ghosh and Caliskan (2023b)’s method. While their method performed resizing of all images, we skip this step since all images across all prompts are generated as 512x512 pixels.

We compute average cosine similarity scores across pairwise comparisons of each image between two image sets. That is to say, for 100 images each generated as the outputs of prompt A to prompt B, we compare each image within the output set of A with each of the 100 images generated for B. This leads to $100 \times 100 = 10000$ comparisons, and we then compute the average score and report that to be the average cosine similarity of the two sets of images.

While CLIP does embed biases (Caliskan, Bryson, and Narayanan 2017), Ghosh and Caliskan (2023b)’s study using CLIP-cosine similarity to compare Stable Diffusion outputs noted that the patterns explicated are indicative of biases embedded within Stable Diffusion, not CLIP. There are two strong reasons for using CLIP-mediated cosine similarity. Firstly, since Stable Diffusion is built upon a CLIP architecture, the CLIP-generated embeddings from its results are likely similar to embeddings used within the process of image formation by Stable Diffusion, and using a different model to produce embeddings from Stable Diffusion outputs would introduce a different set of biases (that of the external model). Secondly, CLIP-mediated cosine similarity is a superior form of the method over other approaches such as RGB (red-green-blue) cosine similarity since the latter only compares raw RGB pixel values across images. While this can be effective for comparing differences in skin tones or background colors, CLIP-generated embeddings encode more information from images. Thus, CLIP-cosine similarity is an effective method for our purpose.

We also supplement results from CLIP-cosine similarity comparisons with manual qualitative verification. Given the subjectivity of manual analysis of images, we attempt to only use this method to point out the presence of objects, colors, background details, and other patterns we believe a vast majority of readers of this paper would agree with.

4 Findings

We report on the various CLIP-cosine similarity comparisons and manual verification results across Caste-Only and Caste-Occupation prompts. As illustrative examples, we provide 2×2 grids of outputs for each prompt, similar to Ghosh and Caliskan (2023b).



Figure 2: Illustrative examples of Stable Diffusion outputs for Caste-Only prompts, in 2x2 grids.

4.1 Findings Across Caste-Only Prompts

To examine how Stable Diffusion outputs represent caste, we compared the 100 image outputs for the prompt ‘Indian person’ with those of the other Caste-Only prompts. We first observed that the images of ‘Indian person’ (Figure 2A) are highly similar to those of ‘Indian high-caste person’ (Figure 2B) with an average CLIP-cosine similarity comparison score of 0.77. Furthermore, ‘Indian person’ is deemed most similar to ‘Indian Brahmin person’ (Figure 2C) and ‘Indian Kshatriya person’ (Figure 2D) with both comparisons yielding average CLIP-cosine similarity scores of 0.76. Closely following is the comparison between ‘Indian person’ and ‘Indian Vaishya person’, (Figure 2E) which produced a CLIP-cosine similarity score of 0.71.

However, ‘Indian person’ is less similar to ‘Indian low-caste person’ (Figure 2F), with a CLIP-cosine similarity score of 0.63. This score is also comparable to that of comparing ‘Indian person’ to ‘Indian Adivasi person’ (Figure 2H), which yields a score of 0.58. This drop can perhaps be attributed to the fact that while the outputs for ‘Indian person’ alongside those for ‘Indian high-caste person’ and prompts with Savarna titles show faces with blurred or neutral backgrounds, those for ‘Indian low-caste person’ and ‘Indian Adivasi person’ feature fields and farmland or mud huts in the backgrounds of most images. It is also notable that while outputs for both ‘Indian person’ and ‘Indian Adivasi person’ show similar colors (such as yellow, saffron, and red), the colors for Adivasi people are represented either as facepaint or in neckwear, conforming to a stereotypical depiction of tribal peoples (Cramer 2005) and possibly explaining the drop in score as compared to ‘Indian low-caste person’. Furthermore, while outputs for the aforementioned prompts mostly show headshots, those for ‘Indian low-caste person’ feature many images with individu-

als from the chest-up and shirtless, further amplifying the stereotype of low-caste and Adivasi people being peasants and farmworkers living in rural areas. A similar rationale is also considered accurate for the fact that CLIP-cosine similarity score comparing the outputs of ‘Indian person’ to ‘Indian Shudra person’ (Figure 2G) is only 0.59). Alongside the aforementioned patterns of increased waist/chest-up images showing shirtless individuals and a prominence of farmland in backgrounds which indicate conditions of rural living, the further drop of this score from the 0.63 reported for ‘Indian low-caste’ person can perhaps be attributed to the fact that 19% of the images in the output of ‘Indian Shudra person’ are in grayscale, though we cannot be certain about this categorically being the reason for such a drop.

The lowest CLIP-cosine similarity score is observed when comparing the outputs of ‘Indian person’ with those of ‘Indian Dalit person’ (Figure 2I, only producing a score of 0.37. Manual examination of the latter set of images produces clear evidence explaining this drop: these images seem to mostly ignore the fact that the prompt asks for a single person and instead heavily feature (in 81% images) large groups of people in a single image. Furthermore, while outputs for ‘Indian person’ feature several aforementioned colors, outputs for ‘Indian Dalit person’ firmly show images containing the color blue, in 81% images. This can perhaps be attributed to the fact that blue was the color of the Independent Labour Party started by Dr. Ambedkar towards the upliftment of caste-oppressed groups and has been used in several Dalit protests in recent history, making the color almost synonymous with Dalit struggle (Rajan and Venkatraman 2017). More recently, blue has been associated with Dalit rights in the electoral victory of Azad Samaj Party president Chandrashekar Azad in the 2024 Indian Lok Sabha elections, and his campaign featured thousands marching in blue attire.



Figure 3: Illustrative examples of Stable Diffusion outputs for Caste-Occupation prompts, in 2x2 grids.

4.2 Findings Across Caste-Occupation Prompts

For the Caste-Occupation prompts, we observe that CLIP-cosine similarity of the outputs for ‘Indian person, at work’ (Figure 3A) and ‘Indian high-caste person, at work’ (Figure 3B) is 0.76. Furthermore, the comparison of the outputs of ‘Indian person, at work’ with those of ‘Indian Brahmin person, at work’ (Figure 3C), ‘Indian Kshatriya person, at work’ (Figure 3D) and ‘Indian Vaishya person, at work’ (Figure 3E) yields the CLIP-cosine similarity scores of 0.74, 0.72, and 0.67 respectively. Manual examination shows a large majority of images across outputs (77-81%) depict individuals seated at a desk or on some form of chair, and either working on a laptop/typing on a keyboard or writing in notebooks as indicators of being ‘at work’. There are differences in the backgrounds; while the backgrounds for ‘Indian person, at work’ (Figure 3A) show office-like backgrounds with other equipment, most of the images for other three prompts depict walls and backgrounds in more rural settings.

However, when looking at prompts representing caste-oppressed groups, we observe a shift in the pattern. Firstly, the result of comparing the outputs of ‘Indian person’ with those of ‘Indian low-caste person, at work’ (Figure 3F), ‘Indian Shudra person, at work’ (Figure 3G), ‘Indian Adivasi person, at work’ (Figure 3G) and ‘Indian Dalit person, at work’ (Figure 3I) yield CLIP-cosine similarity scores of 0.53, 0.47, 0.46 and 0.44, respectively. This represents a sharp drop in similarity scores from the aforementioned scores in the 0.74-0.68 range, and to explain these, we once again turn to manual qualitative analysis of images.

Across these outputs, a strikingly observable point of difference with the outputs for ‘Indian person, at work’ (Figure 3A) is that *not a single person* is shown to be working with laptops/keyboards or writing into notebooks. Rather, the interpretation of ‘at work’ for these outputs (Figures 3F,

3G, 3H, and 3I) primarily represents actions and occupations such as spinning cloth at a wheel, laying bricks, stacking large sacks or logs of wood, working on farmland, mixing cement, carrying jugs of water or other items, and sweeping streets, identified through the presence of objects associated with those actions/occupations (such as bricks, logs, sacks, etc.) present within images. Such objects are seen with comparable prominence across outputs for each prompt.

We also note that the outputs for all of these prompts prominently feature rural-appearing backgrounds such as farmlands, mud-colored floors and walls, or sandy streets, as illustrated in Figures 3F, 3H, and 3I. Perhaps the only exception could be the outputs for ‘Indian Shudra person, at work’ (Figure 3G), but that could also be because of the common occurrence of grayscale images (27%) within this set, in keeping with a similar pattern of featuring grayscale images to represent Shudras. It is also evident that the addition of the phrase ‘at work’ for prompts around Adivasis or Dalits still maintained similar color palettes, with outputs for Adivasi people (Figures 2H and 3H) continuing to show lots of bright colors and those for Dalits (Figures 2I and 3I) still showing a prominent blue pattern. However, in the context of the general pattern of these images showing a prominence of mud-colored or brown backgrounds, these splashes of color seem to stand out even more.

Finally, it is important to highlight that the images generated for ‘Indian Dalit person, at work’ (Figure 3I) were different from those for ‘Indian Dalit person’ (Figure 2I), in the way that the former succeeded in showing a single person per image where the latter did not. It is not immediately clear why that is the case, and is probably unfair to assume that the addition of the phrase ‘at work’ is the cause of this difference, but it is nevertheless interesting to point out.

5 Analysis: ‘Castelessness’ in Stable Diffusion, but with a Few Harmful Stereotypical Exceptions

Through CLIP-cosine similarity comparisons and manual evaluation of Stable Diffusion outputs across Caste-Only and Caste-Occupation prompts, we observe patterns of Stable Diffusion embedding an apparent ‘castelessness’ within prompts containing Savarna identities and harmful stereotypes around those with caste-oppressed ones.

To understand how Stable Diffusion interprets caste, we document the high CLIP-cosine similarity scores obtained by comparing the outputs of ‘Indian person’ (Figure 2A) with those of ‘Indian high-caste person’ (Figure 2B), ‘Indian Brahmin person’ (Figure 2C), ‘Indian Kshatriya person’ (Figure 2D), and ‘Indian Vaishya person’ (Figure 2E) all falling in the 0.77-0.71 range. In contrast, the outputs for ‘Indian low-caste person’ (Figure 2F), ‘Indian Shudra person’ (Figure 2G), ‘Indian Adivasi person’ (Figure 2H) and ‘Indian Dalit person’ (Figure 2I) only receive CLIP-cosine similarity scores of 0.63, 0.59, 0.58 and 0.37 respectively, in comparison with ‘Indian person’. Furthermore, even for Caste-Occupation prompts, CLIP-cosine similarity scores comparing ‘Indian person, at work’ with outputs for prompts with Savarna labels, i.e. ‘high-caste’ (Figure 3B), ‘Brahmin’ (Figure 3C), ‘Kshatriya’ (Figure 3D), and ‘Vaishya’ (Figure 3E) produce scores in the 0.76-0.67 range, whereas comparisons to prompts with caste-oppressed labels – ‘low-caste’ (Figure 3F), ‘Shudra’ (Figure 3G), ‘Adivasi’ (Figure 3H) and ‘Dalit’ (Figure 3I) – show scores in the 0.53-0.44 range. It is important to recognize that these gaps represent significant differences in pairwise comparisons, especially given that each CLIP-cosine similarity score is an average of 10000 individual comparisons. Collectively, these results are indicative of a pattern that Stable Diffusion considers the default ‘Indian person’ to be of high-caste, and working in some form of office-based work involving laptops/keyboards. Such a pattern is problematic on a few different levels, beginning with the fact that this result is misrepresentative of the Indian population, since high-caste individuals were found to only be 48% of the Indian population according to the last census (Chandramouli 2011), which implies that the default or ‘average’ Indian person is more likely to be lower-caste than high. Furthermore, less than 42% of India’s working population is employed at what is considered a ‘non-household and non-agricultural occupation’ of which white collar jobs form less than a quarter, whereas agricultural occupations form over 51% (Chandramouli 2011), implying that the average ‘Indian person, at work’ is likely not working with a laptop/computer.

Perhaps more dangerously, the presence of such a pattern might not present as a problem to researchers within India or in the Indian diaspora, or motivate them to address it. Even though only 48% of India’s population is Savarna, they historically and currently do have access to more resources and opportunities to pursue higher education than caste-oppressed individuals. It is probably a fair assumption that the average Indian person regularly interacting with or studying/researching T2I tools such as Stable

Diffusion are more likely to be higher-caste than low. Such a Savarna researcher might therefore find themselves represented within Stable Diffusion outputs for ‘Indian person’, or at the very least agree with such outputs being closer to those for prompts with Savarna labels. Such results create a sense of *castelessness*: a “rendering in which Savarna individuals are able to frame themselves as largely casteless (and meritorious), while lower-caste individuals are seen as still marked by caste.” (Vaghela, Jackson, and Sengers 2022) This also comes at a time when, both in India and the Indian diaspora, individuals and communities occupying privileged positions in terms of caste and religion are spreading moral panic about how their ways of life are being threatened by ‘Others’. Similar to waves of moral panic within the US that White Americans are being discriminated against (Reed 2018), organizations such as the Hindu American Foundation and the Coalition of Hindus of North America, made up mostly of Savarna individuals, opposed legislation across the US around extending affirmative action to caste, calling such protections unnecessary and unconstitutional. In India, such a moral panic often manifests through criticisms of caste-based reservation systems, arguing that Savarna individuals are more meritorious but have access to fewer opportunities due to such a system. The growing usage of T2Is in India and the diaspora coupled with the fact that Savarna individuals control the resources necessary to meaningfully dismantle castelessness within T2I outputs thus lead to the likely scenario that such an alarming pattern will be allowed to persist.

When considering this pattern of castelessness where only the “lower-caste individuals are seen as still marked by caste” (Vaghela, Jackson, and Sengers 2022), it is also important to examine *how* the marking occurs. Harkening back to the low similarity scores for prompts embedding caste-oppressed labels in comparison to ‘Indian person’ (0.63-0.37) and ‘Indian person, at work’ (0.53-0.44), these low scores can be attributed to the presence of visible differences, which function as markers of caste-oppression. When Stable Diffusion outputs depict caste-oppressed individuals working in what can only be described as unspecified blue-collar work involving manual labor (Section 4.2) typically on farmlands and the worker being shirtless (Section 4.1), the model demonstrates the embedding of the stereotype that caste-oppressed individuals are impoverished, live in rural areas, and only perform agricultural or manual labor. It might seem, following the aforementioned comparison to population statistics, that these results are accurate: indeed, as a large majority of caste-oppressed people in India lives in rural areas and practices agriculture, such depictions might appear to be representative of the community, over and above the fact that prior work has shown representations of India/South Asia to generally skew impoverished (Qadri et al. 2023). However, the fatal flaw in these results is the fact that *not a single depicted individual* across a total of 800 images spanning 8 prompts with caste-oppressed labels – 4 each for Caste-Only and Caste-Oppressed containing ‘low-caste’, ‘Shudra’, ‘Adivasi’, and ‘Dalit’ – seems to deviate from this pattern. Across 800 images, there is not a single representation of a caste-oppressed individual with visible markers of “traditional” forms of success and

affluence. Indeed, there is not even a single depiction of a caste-oppressed person wearing a collared shirt, one of the oldest markers of wealth and prosperity across societies (Brough 2013). The pattern of impoverished depictions of caste-oppressed individuals is thus not Stable Diffusion accurately depicting a group based on statistical information, but rather a pervasive pattern of the representational harms of stereotyping (casting caste-oppressed groups as poor) and erasure (not showing a single instance of a caste-oppressed individual being wealthy or working a white collar job). As such models reach a global user base, most of whom might not have a nuanced understanding of caste, the propagation of this stereotype can amplify such harms at scale.

The strongest of such representational harms is seen within depictions of Dalit individuals, the most historically marginalized of all the groups studied here. Stable Diffusion outputs around Dalits show the overwhelming pattern of Dalits organizing in large groups at protests (Figure 2I) or performing menial labor (Figure 3I), as this set of outputs receives a lowly 0.37 CLIP-cosine similarity score when compared to the outputs of ‘Indian person’. Over 10000 comparisons, such a low score is indicative of a large mismatch between ‘Indian person’ and ‘Indian Dalit person’, echoing previous work showing the large distancing between T2I depictions of default/privileged and marginalized identities (Ghosh and Caliskan 2023b), and the manually-observed details of this mismatch are alarming. Of particular note is the observation that 81% of the 100 images for ‘Indian Dalit person’ shows the presence of multiple people. This is in complete opposition to the fact that the full prompt (see Section 3.1) explicitly asks for ‘a front-facing photo of a Dalit person from India’, and that a similar construction for every one of the other 17 prompts (and 1700 images) yields one individual per image in almost all cases. While it can be argued that the group/rally-style images are results of the recent rise in such images on the Internet of Dalit people marching and rallying to campaign for Azad Samaj Party president Chandrashekar Azad in the 2024 Indian Lok Sabha election, the strong stereotyping of Dalits as protesters and at rallies is problematic for two reasons. Firstly, such a depiction causes representational harms such as dehumanization by implying that a Dalit person exists, by default, within a protest or rallying space and not as their own individual, which leads to a loss of individual agency of Dalit people. Secondly, it fuels an anger that has been brewing in India over the past decade, primarily within privileged individuals, that traditionally marginalized people such as Dalits are protesting “too much” (Kumbhat 2016). Such sentiments are not confined to India, as the New York Times faced backlash for introducing an article on Dalit protests in India with the line “But today there are Dalit millionaires, so why are they protesting?” (Sengar 2018) The fact that Stable Diffusion outputs predominantly portrays Dalits in large groups can cause serious harm to them.

It is clear from our work that Stable Diffusion has an imperfect and incomplete understanding of caste, and we advocate for a change in this vein.

6 Towards Equitable and Accurate Representation of Caste within T2Is

In this section, we present design recommendations towards obtaining a more fine-grained representation of caste within T2I outputs. These recommendations align with and extend other research (e.g., Gadiraju et al. 2023; Ghosh and Caliskan 2023a; Mack et al. 2022; Qadri et al. 2023) advocating for community- and human-centered approaches towards redesigning the LLMs underpinning T2Is.

In similar studies explicating representational harms within outputs of T2Is or Generative AI (GAI) tools, as suffered by specific communities because of their historic marginalization due to a shared identity (e.g., Gadiraju et al. 2023; Mack et al. 2022; Qadri et al. 2023), a common design recommendation is for stronger community involvement. Such recommendations broadly revolve around soliciting stronger community participation in T2I design, through processes such as community-centered data collection and annotation procedures. These recommendations, either implicitly or explicitly, assume that some or all community members would be willing and able to participate in such efforts. For caste-oppressed individuals, this assumption might not hold true for a majority of the community (Shah 2016).

Firstly, a significant portion of caste-oppressed communities within India might not have any interaction with T2Is that they are aware of. Using and working with GAI tools and T2Is requires a level of access to technology and infrastructure such as reliable high-speed Internet connections and devices with high-quality graphics cards to be able to effectively render outputs, as well as proficiency around prompting techniques, which is likely not commonly made available to caste-oppressed communities. As a result, the sort of data that researchers might imagine gathering might not be available or even feasible (Sambasivan 2021). The most likely members of caste-oppressed communities who might be willing and able to participate in community-centered calls towards better representation of caste-oppressed groups are individuals who have acquired traditional metrics of success such as higher education and relative affluence, and possibly the infrastructure and skills to be interacting with T2Is. While their input will undoubtedly be valuable, it is also unfair to ask them to speak for their entire community. This ask will impose a burden to be perfect spokespeople for all caste-oppressed communities, subjecting them to ‘minority tax’ (Rodríguez et al. 2021) by being tokenized as model representatives of their community.

Furthermore, even if the question of access to knowledge and technology is bypassed, it is important to consider the politics of data collection from caste-oppressed individuals. Within the Indian diaspora or in regions where individuals of multiple castes coexist, there might be serious pushback against providing data that makes caste-oppressed individuals reveal their identity for fear of suffering caste-based violence and discrimination (Zwick-Maitreyi et al. 2018). As an aspect of identity around which hate crimes occur, caste might not be something that individuals are willing to disclose, even if for lofty goals such as algorithmic fairness (Andrus and Villeneuve 2022). Community-centered data

collection and annotation procedures with caste-oppressed individuals where they identify or label ‘accurate’ depictions of their identities might also be problematic. The term ‘caste-oppressed’ encompasses millions of people across large parts of India and the world, spanning various castes and tribes. There thus cannot be a single idea of what a ‘caste-oppressed’ person look like, or even what a Dalit or an Adivasi person looks like, and working towards creating such a representation might do more harm than good (Atri 2019). In collecting data from a vulnerable community like caste-oppressed groups, researchers must practice cultural sensitivity, work with local community experts to not assume researcher perspectives as knowing what is ‘good’ or ‘fair’ for the community (Nathan et al. 2017), and always offer right of refusal (Ghosh and Chatterjee 2024).

Rather, one of the strongest ways towards equitable representation of caste-oppressed communities is to raise awareness about the existence and proliferation of GAI tools, and ‘humbly build grassroots commitments’ (Sambasivan et al. 2021). Such an education must begin from the very beginning of what GAI tools and T2Is are, even by sharing simplified versions of definitions that do not go beyond the concept of such tools taking in prompts and outputting text/images. Furthermore, especially for individuals who believe that they would never use or have use for such tools, it is important to inform them that they may or may not already be influenced by the outputs of such tools (Thakkar, Kumar, and Sambasivan 2020). For instance, in the ongoing 2024 Indian General elections, GAI tools are being used to generate campaign materials where deepfake videos of candidates are being made that address voters by name (Raj 2024a). Education around T2Is must also include conversations around trust, since individuals unfamiliar with T2Is might not have reason to suspect a video as deepfaked (Kapania et al. 2022).

Another important step is to raise global awareness within researchers interested in fairness/harm reduction in T2Is or machine learning, about the fact that caste-based harms can happen through such outputs and is an important problem to focus on. Caste has been an aspect of identity historically sidelined in progressive and feminist movements (Rege 1998) and continues to remain an identity around which fairness research is not happening. Indeed, as Sambasivan et al. (2021) notes, a Western orientation around fairness in machine learning fails to even consider caste as an axis along which discrimination occurs. Through (hopefully) the propagation of this work and the undertaking of other research, we invite the spreading of a larger awareness around caste-based harms in GAI tools. Our work is by no means all-conclusive around casteist representations in T2Is, and perhaps one of the most logical extensions of it could be a human subjects study. One such open question is how T2I outputs being operationalized into downstream tasks such as generating marketing or campaign content using depictions like the aforementioned for Indian people or people of a certain caste risk can cause allocational or representational harms (Barocas et al. 2017) upon traditionally caste-marginalized people, as documented through human subjects studies. We invite researchers, especially those familiar with caste as an identity and reflexive about their own posi-

tions within such systems, to undertake such work and contribute to the global research community. In this vein, we appreciate the work done in spotlighting how caste is represented in machine learning tools by researchers such as Qadri et al. (2023) and Dev et al. (2024), among others.

It is also important to define what ‘equitable representation’ for caste-oppressed communities within T2Is might even look like. Prevalent concepts of equity within ML are often Western-centric and should not be applied to Indian contexts directly, as there might be cultural mismatches (Sambasivan et al. 2020, 2021). It is important to consider local contexts such as caste-based reservations, and what that might mean for the conception of ‘fairness’. Furthermore, especially as popular models might not yet have a fine-grained understanding of caste, there is a scenario where working towards fairness by adding more data actually does more harm than good. For instance, given that a person’s caste can be ascertained from their last name, adding demographic information to train models might then create systems that are capable of labeling caste, which may in turn be used to discriminate against caste-oppressed groups in downstream tasks such as resume screening. Especially in the context of caste, deliberate confusion of models might actually be a good thing (Sambasivan et al. 2021).

Above all, it is critical to advocate for social change and the prevention of caste-based oppression. It is a well-established notion within research around machine learning tools such as T2Is that they only amplify and embed the biases which are already within society. At their heart, T2Is are agents of social power, and the way to properly combat the propagation of negative stereotypes is to rise up to dismantle the oppressive systems of power that create them. Therefore, efforts into achieving equitable representations of caste-oppressed groups must also include meaningful support towards their real-life causes, providing material support and using individual positions of privilege – caste-based or not – to stand up against unjust caste-oppressive practices.

7 Limitations and Conclusion

A limitations of this study is that it has been led by an researcher identifying as caste-privileged. While this does not invalidate the work, it leaves open the possibility that researcher with stronger first-hand epistemic experiences around caste-based marginalizations would have deeper insights into the interpretations and representations of caste within these findings. We invite further explorations by researchers who have been marginalized by virtue of their castes, and are currently seeking such collaborations.

In conclusion, this paper highlights representations, interpretations, and stereotypes around caste within Stable Diffusion outputs, and how they perpetuate systems of ‘castelessness’ prevalent within Indian society and the diaspora. We provide design recommendations towards equitable and just representations of caste-oppressed identities, with the hope that this work encourages further focus around advocating for equitable representations beyond non-Western contexts.

Researcher Positionality Statement

The first author of this paper identifies as an individual born and raised in the Bengal region of India, and is currently an international student in the United States. Their caste identity is Kulin Kayastha, which is a caste-privileged group from the region. Though the author is relatively caste-privileged, they have still personally experienced casteist comments from individuals of higher castes than them.

Ethical Statement

In doing any research with T2Ist that involve the sharing of sample generated outputs alongside their associated prompt, there is always the possibility of an adverse impact: by putting an image-caption combination on the Internet, researchers might unwittingly be providing more data to be scraped and used for model training in the future. This is especially impactful for images such as the ones we present, because of the stereotyping within them. Therefore, pursuant to Ghosh and Caliskan (2023b), we too will provide blurred images upon publication, both for the images in the paper and those in the dataset.

References

- Abia, O.; Amalu, N. S.; and Ariche, C. K. 2021. Osu caste system and human rights in Igboland, 1900-2017. *Global Journal of Social Sciences*, 20(1): 69–76.
- Andrus, M.; and Villeneuve, S. 2022. Demographic-reliant algorithmic fairness: Characterizing the risks of demographic data collection in the pursuit of fairness. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1709–1721.
- Attri, P. 2019. ‘You don’t look Dalit’ and other things ‘upper castes’ must stop saying to Dalits immediately. *The Print*.
- Barocas, S.; Crawford, K.; Shapiro, A.; and Wallach, H. 2017. The problem with bias: Allocative versus representational harms in machine learning. In *9th Annual conference of the special interest group for computing, information and society*.
- Benjamin, R. 2019. *Race after technology: Abolitionist tools for the new Jim code*. John Wiley & Sons.
- Bharath, D. 2023. What is India’s caste system? Is it contentious in U.S.? *The Associated Press*.
- Bianchi, F.; Kalluri, P.; Durmus, E.; Ladhak, F.; Cheng, M.; Nozza, D.; Hashimoto, T.; Jurafsky, D.; Zou, J.; and Caliskan, A. 2023. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 1493–1504.
- Boon, J. A. 1977. *The anthropological romance of Bali 1597-1972: dynamic perspectives in marriage and caste, politics and religion*, volume 1. CUP Archive.
- Bowker, G.; and Star, S. L. 1999. Sorting things out. *Classification and its consequences*, 4.
- Brough, D. 2013. The classic white formal shirt: a powerful emblem of social change. In *Proceedings of the 15th Annual International Foundation of Fashion Technology Institutes Conference*, 1–14. International Foundation of Fashion Technology Institutes-IFFTI.
- Caliskan, A.; Bryson, J. J.; and Narayanan, A. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334): 183–186.
- Chandramouli, C. 2011. 2011 Census Report. Technical report, Government of India.
- Collins, P. H. 1990. Black feminist thought in the matrix of domination. *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*, 138(1990): 221–238.
- Cramer, R. A. 2005. *Cash, color, and colonialism: The politics of tribal acknowledgment*. University of Oklahoma Press.
- Crenshaw, K. 2017. *On intersectionality: Essential writings*. The New Press.
- Dammu, P. P. S.; Jung, H.; Singh, A.; Choudhury, M.; and Mitra, T. 2024. “They are uncultured”: Unveiling Covert Harms and Social Threats in LLM Generated Conversations. *arXiv preprint arXiv:2405.05378*.
- Dev, S.; Goyal, J.; Tewari, D.; Dave, S.; and Prabhakaran, V. 2024. Building Socio-culturally Inclusive Stereotype Resources with Community Engagement. *Advances in Neural Information Processing Systems*, 36.
- Dev, S.; Li, T.; Phillips, J. M.; and Srikumar, V. 2020. On measuring and mitigating biased inferences of word embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 7659–7666.
- Du, C.; Li, Y.; Qiu, Z.; and Xu, C. 2024. Stable diffusion is unstable. *Advances in Neural Information Processing Systems*, 36.
- Gadiraju, V.; Kane, S.; Dev, S.; Taylor, A.; Wang, D.; Denton, E.; and Brewer, R. 2023. “I wouldn’t say offensive but...”: Disability-Centered Perspectives on Large Language Models. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 205–216.
- Gautam, S.; Venkit, P. N.; and Ghosh, S. 2024. From melting pots to misrepresentations: Exploring harms in generative ai. *arXiv preprint arXiv:2403.10776*.
- Ghosh, S.; and Caliskan, A. 2023a. ChatGPT Perpetuates Gender Bias in Machine Translation and Ignores Non-Gendered Pronouns: Findings across Bengali and Five other Low-Resource Languages. *AAAI/ACM Conference on AI, Ethics, and Society 2023*, 901–912.
- Ghosh, S.; and Caliskan, A. 2023b. ‘Person’ == Light-skinned, Western Man, and Sexualization of Women of Color: Stereotypes in Stable Diffusion. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*, 6971–6985. Singapore: Association for Computational Linguistics.
- Ghosh, S.; and Chatterjee, S. 2024. Misgendering and Assuming Gender in Machine Translation when Working with Low-Resource Languages. In Monzo-Nebot, E.; and Tasa-Fuster, V., eds., *Gendered Technology in Translation and Interpreting: Centering Rights in the Development of Language Technology*. Routledge.

- Ghosh, S.; Venkit, P. N.; Gautam, S.; Wilson, S.; and Caliskan, A. 2024. Do Generative AI Models Output Harm while Representing Non-Western Cultures: Evidence from A Community-Centered Approach. *arXiv preprint arXiv:2407.14779*.
- Hall, S. 1997. *Representation: Cultural representations and signifying practices*, volume 2. Sage.
- Harad, T. 2023. After US, Ambedkarites in UK Renew Their Fight for Caste Legislation. *The LiveMint*.
- Ho, P. 2003. Chinese Society: Change, Conflict and Resistance.
- Kapania, S.; Siy, O.; Clapper, G.; SP, A. M.; and Sambasivan, N. 2022. "Because AI is 100% right and safe": User attitudes and sources of AI authority in India. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Kapoor, S. 2003. BR Ambedkar, WEB DuBois and the process of liberation. *Economic and Political Weekly*, 5344–5349.
- Kaur, H. 2023. Seattle becomes the first city in the US to ban caste discrimination. *CNN*.
- Kumbhat, U. 2016. The Country Doth Protest Too Much. *Chicago Maroon*.
- Lawlor, P.; and Chang, J. 2023. The Generative AI economy: Worth up to 7.9T.
- Luccioni, A. S.; Akiki, C.; Mitchell, M.; and Jernite, Y. 2023. Stable Bias: Analyzing Societal Representations in Diffusion Models. *NeurIPS Datasets and Benchmarks*.
- Mack, A.; Qadri, R.; Denton, R.; Kane, S. K.; and Bennett, C. L. 2024. "They only care to show us the wheelchair": Disability Representation in text-to-image AI models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*.
- Mack, K. A.; McDonnell, E.; Potluri, V.; Xu, M.; Zabala, J.; Bigham, J.; Mankoff, J.; and Bennett, C. 2022. Anticipate and adjust: Cultivating access in human-centered methods. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–18.
- Morgan, K. P. 2018. Describing the Emperor's New Clothes: Three myths of educational (in-)equity. In *The gender question in education*, 105–122. Routledge.
- Nathan, L. P.; Kaczmarek, M.; Castor, M.; Cheng, S.; and Mann, R. 2017. Good for whom? Unsettling research practice. In *Proceedings of the 8th International Conference on Communities and Technologies*, 290–297.
- OpenAI. 2021. CLIP: Connecting text and images.
- Qadri, R.; Shelby, R.; Bennett, C. L.; and Denton, E. 2023. AI's Regimes of Representation: A Community-centered Study of Text-to-Image Models in South Asia. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 506–517.
- R, J. 2024. Karnataka government does a U-turn, to continue financial assistance to SC-ST students to pursue PhD in foreign universities under Prabuddha scheme. *The Hindu Impact*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Raj, S. 2024a. How A.I. Tools Could Change India's Elections. *The New York Times*.
- Raj, S. 2024b. 'Chamar-Dalit quit India', 'Brahmin-RSS zindabad' slogans found on JNU hostel walls, students express anger. *Maktoob Media*.
- Rajan, B.; and Venkatraman, S. 2017. Fabric-rendered identity: A study of Dalit representation in Pa. Ranjith's Atakathi, Madras and Kabali. *Artha Journal of Social Sciences*, 16(3): 17–37.
- Rajvanshi, A. 2023. How Seattle Became the First American City to Ban Caste Discrimination. *TIME Magazine*.
- Reed, A. 2018. Knowing Our Demons. *The Immanent Frame*.
- Rege, S. 1998. Dalit women talk differently: A critique of 'difference' and towards a Dalit feminist standpoint position. *Economic and Political Weekly*, WS39–WS46.
- Rodríguez, J. E.; Wusu, M. H.; Anim, T.; Allen, K.-C.; and Washington, J. C. 2021. Abolish the minority woman tax! *Journal of Women's Health*, 30(7): 914–915.
- Saha, S. 2020. Caste, Materiality and Embodiment. *CASTE: A Global Journal on Social Exclusion*, 1(1): 200–216.
- Sambasivan, N. 2021. Seeing like a dataset from the global south. *Interactions*, 28(4): 76–78.
- Sambasivan, N.; Arnesen, E.; Hutchinson, B.; Doshi, T.; and Prabhakaran, V. 2021. Re-imagining algorithmic fairness in india and beyond. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, 315–328.
- Sambasivan, N.; Arnesen, E.; Hutchinson, B.; and Prabhakaran, V. 2020. Non-portability of algorithmic fairness in india. *arXiv preprint arXiv:2012.03659*.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. LAION-5B: An Open Large-Scale Dataset for Training Next Generation Image-Text Models. *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS 2022)*.
- Sengar, S. 2018. NYT's Tweet 'Today There Are Dalit Millionaires... Why Are They Protesting?' Reflects Its Myopic Understanding Of India. *Scroll.in*.
- Shah, S. 2016. MissionCashless: Few use mobiles, fewer know what internet is in adivasi belts of Madhya Pradesh. *Scroll.in*.
- Shantha, S. 2024. From Segregation to Labour, Manu's Caste Law Governs the Indian Prison System. *The Wire*.
- Shelby, R.; Rismani, S.; Henne, K.; Moon, A.; Ros-tamzadeh, N.; Nicholas, P.; Yilla, N.; Gallegos, J.; Smart, A.; Garcia, E.; et al. 2022. Sociotechnical harms: Scoping a taxonomy for harm reduction. *arXiv*.
- Shetty, V. R. 1979. *Apartheid in India: A Document for Foreigners*. Dalit Action Committee: Bangalore.

Singhal, A.; et al. 2001. Modern Information Retrieval: A Brief Overview. *IEEE Data Engineering Bulletin*, 24(4): 35–43.

StabilityAI. 2022. Stable Diffusion Launch Announcement. *Stability AI Blog*. Has a quote there about 10000 beta testers, but does not say what parameters and what types of images.

Thakkar, D.; Kumar, N.; and Sambasivan, N. 2020. Towards an AI-powered future that works for vocational workers. In *proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13.

Tiwari, P.; Kumar, A. C.; Chandrabose, A.; et al. 2022. Casteism in India, but not racism—a study of bias in word embeddings of Indian languages. In *Proceedings of the First Workshop on Language Technology and Resources for a Fair, Inclusive, and Safe Society within the 13th Language Resources and Evaluation Conference*, 1–7.

Vaghela, P.; Jackson, S. J.; and Sengers, P. 2022. Interrupting merit, subverting legibility: Navigating caste in ‘casteless’ worlds of computing. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, 1–20.

Wilkerson, I. 2020. *Caste: The origins of our discontents*. Random House.

Wolfe, R.; Yang, Y.; Howe, B.; and Caliskan, A. 2023. Contrastive Language-Vision AI Models Pretrained on Web-Scraped Multimodal Data Exhibit Sexual Objectification Bias. *The ACM Conference on Fairness, Accountability, and Transparency*.

Zwick-Maitreyi, M.; Soundararajan, T.; Dar, N.; Bheel, R.; and Balakrishnan, P. 2018. Caste in the United States. A Survey of Caste among South Asian Americans. USA.