ISSN: 1304-4796
Manisa Celal Bayar Üniversitesi Sosyal Bilimler Dergisi
Manisa Celal Bayar University Journal of Social Sciences
2026; 24 (2);800-822

Reimagining Leadership: Visual Biases in AI-Generated Portraits

Taner Kızılhan[1, 2], Betül Gökçin[3]

ABSTRACT

This study examines how generative AI systems reproduce bias in leadership representations. Despite the use of neutral prompts, analysis of 71 images produced by eight different models revealed clear patterns of bias related to gender, ethnicity, and age. Strong leader figures were predominantly represented as young, Western, and male, while women and ethnic minorities were more often depicted as passive or "visionless." Findings suggest that generative AI systems do not simply reflect data, they replicate and reinforce cultural norms. The research emphasizes the need to move beyond evaluating AI based solely on technical performance, calling instead for assessments based on representational justice. The study argues that generative AI has the capacity to subtly reproduce existing social biases. To ensure diversity and equity in representation, AI training datasets must be developed with inclusive and fair principles. Otherwise, today's biases risk becoming tomorrow's digital myths through the uncritical deployment of such technologies.

Keywords: *Algorithmic Bias, Gender Representation, AI and Visual Culture, Feminist Technology Critique, Symbolic Violence*

Liderliği Yeniden Düşünmek: Yapay Zeka Tarafından Üretilen Portrelerde Görsel Yanlılık

ÖZET

Bu çalışma, üretken yapay zeka sistemlerinin liderlik temsillerinde önyargıları nasıl ürettiğini ve yeniden ürettiğini ele almıştır. Tarafsız biçimde hazırlanan istemlere rağmen, kullanılan sekiz farklı modelden elde edilen 71 görselde toplumsal cinsiyet, etnik köken ve yaş temelli belirgin bir yanlılık gözlemlenmiştir. Güçlü lider figürleri çoğunlukla genç, Batılı ve erkek karakterlerle temsil edilirken; kadınlar ve etnik azınlıklar daha az belirgin, edilgen ya da "vizyonsuz" figürler olarak yansıtılmıştır. Bu durum, yapay zekâ sistemlerinin yalnızca veriyi değil, kültürel normları da yeniden ürettiğini göstermektedir. Araştırma, teknik başarı ölçütlerinin ötesine geçilerek temsilde adalet ilkesi doğrultusunda değerlendirmeler yapılması gerektiğini vurgulamaktadır. Sonuç olarak,

[1] Contact: taner.kizilhan@adu.edu.tr
[2] Asst. Prof. Dr., Aydın Adnan Menderes Üniversitesi, İletişim Fakültesi, Görsel İletişim Tasarımı Bölümü, Efeler, Aydın, Türkiye, ORCID: 0000-0002-3501-415X.
[3] Undergraduate Student, Aydın Adnan Menderes Üniversitesi, İletişim Fakültesi, Görsel İletişim Tasarımı Bölümü, Efeler, Aydın, Türkiye, ORCID: 0009-0005-8506-2642.

(Submission: 06.07.2025 / Published:30.06.2026)
Doi Number: 10.18026/cbayarsos.1736257
Type of Research: Research Article

üretken yapay zekâlar kültürel önyargıları gizli biçimde pekiştirme potansiyeli taşımaktadır. Temsil çeşitliliğini sağlamak için, yapay zekâ modellerinin eğitim verilerinde eşitlik ve kapsayıcılık ilkeleri gözetilmeli; aksi halde bugünün önyargıları, yarının dijital mitlerine dönüşebilir.

Anahtar Kelimeler: *Algoritmik Yanlılık, Toplumsal Cinsiyet Temsili, Yapay Zeka ve Görsel Kültür, Feminist Teknoloji Eleştirisi, Sembolik Şiddet*

1. INTRODUCTION

In recent years, generative AI models have emerged as groundbreaking tools for creating content. Systems such as DALL-E, Midjourney and Stable Diffusion can generate hyper-realistic visuals from textual prompts alone, producing images that did not previously exist. However, beneath these striking capabilities lies a growing concern about the objectivity and neutrality of such systems. Preliminary investigations suggest that generative AI models often reproduce, and in some cases amplify, social biases embedded within their training data (Tiku et al., 2024). There is an increasing body of evidence pointing to the reinforcement of gender and ethnic stereotypes in AI-generated imagery, prompting critical debates at the intersection of AI ethics, gender representation, and visual culture.

The potential implications of such biases are far from negligible. The ubiquity of AI-produced imagery is increasing, and it is important to consider the potential impact of this on the perpetuation of social stereotypes. Visual content has been shown to have a significant impact on the development of shared imaginaries, and it is therefore possible that the growing use of AI-generated imagery may contribute to the solidification of pre-existing social norms. A recent analysis by The Washington Post found that generative image models often reproduce 'troubling stereotypes'; for example, people in leadership roles are mostly male (Tiku et al., 2024). Similarly, a study by the Brookings Institution found that DALL-E generated twenty images in response to the prompt 'a successful person', all of which depicted young, white male business professionals, despite the prompt not specifying any gender or ethnicity (Baum & Villasenor, 2024). This consistent tendency towards these characteristics suggests that such models implicitly internalize patriarchal and Western-centric conceptions of leadership.

This emerging picture calls into question the “neutrality” often ascribed to AI systems. These algorithms tend to replicate the dominant patterns of representation embedded within the vast datasets on which they are trained. For instance, a 2024 UNESCO report reveals a clear gender bias in popular large language models: 'Women are associated with terms such as "home", "family", and "children" four times more frequently than men, whereas male names are predominantly linked to words such as "job", "executive", "salary", and "career"', (UNESCO, 2024). These ingrained biases in AI systems raise significant ethical concerns at the intersection of technology and society. Yet these issues remain insufficiently explored within the domain of image generation. While scholarships on AI and social bias is expanding, analyses of bias in the visual outputs of generative AI are notably scarce from a social science perspective. Consequently, critical questions, such as whether generative

AI models perpetuate traditional biases related to gender and phenotype and, if so, the mechanisms through which they do so, remain largely unresolved.

2. LITERATURE REVIEW

2.1. Definition and Nature of Artificial Intelligence

'The term "artificial intelligence" was coined in 1956 at a scientific conference held at Dartmouth College in Hanover, New Hampshire' (Özkıran, 2022, p. 1). Despite its long history, there is still no universally accepted definition of the term, largely due to differing disciplinary perspectives. As a result, defining artificial intelligence remains a complex task. To provide a clearer understanding of its nature, Table 1 presents a selection of definitions that have been developed over the years.

Table 1. Definitions of Artificial Intelligence

Source	Definition
NcCarthy et al, 1955	The science and engineering of making intelligent machines (cited in Kolata, 1982): what is really needed as machines that can solve problems, not machines that think as people.
Winston, 1984	The study of ideas that enable computers to be intelligent. (p. 3)
Charniack & McDermott,1985	The study of mental faculties with the use of computational models. (p. 6)
Rich & Knight, 1991	The Study of how to make computers do things, which at the moment, people do better. (p. 3)
Poole, Macworth & Goebel, 1998	The intelligence exhibited by machines software.
Luger & Stubblefield,2008	The branch of computer science that is concerned with the automation of intelligent behavior.
Russell & Norvig, 2010	The study and design of rational agent is a system that perceives its environment and takes actions that maximize its chance of success. Russell & Norvig classify previous definitions of AI into two-by-two matrix:Thinking versus acting; and humanly versus rationality.
Oxford dictionaries, 2014	The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision making, and translation between languages.
AAAI, 2014	Mechanisms underlying thought and intelligent behavior and their embodiment in machines.

Kaynak: El-Hadi, M. (2023). Artificial intelligence background, definitions, challenges and benefits. *CompuNet 31.* Retrieved June 17, 2025 from article_297957_18c63823bf45bdd85ffd54bf0dafa8f5.pdf

This compilation of definitions reflects the evolving perspectives on artificial intelligence alongside the historical development of the field. Early definitions emphasized the problem-solving capabilities of machines (McCarthy et al., 1955) and information processing theories (Nilsson, 1980). Over time, however, the focus shifted towards concepts such as human-like reasoning (Charniak & McDermott, 1985) and rational behavior (Russell & Norvig, 2010). This diversity indicates that AI is grounded in both engineering-oriented applications and theoretical approaches that intersect with cognitive science. More recent definitions emphasize multidimensional features such as the imitation of human intelligence, learning and adaptive decision-making. The definition adopted in this study, due to its clarity and inclusiveness, is that AI "is a technology that enables machines to imitate various complex human skills" (Russell & Norvig, 2020, p. 15).

'Over the past few decades, AI has evolved from a theoretical concept into an integral and pervasive part of our daily lives' (Min, 2023, p. 2). People now use a wide range of AI tools intensively, not only in professional contexts, but also in personal spheres such as learning, production and consumption. One of the most significant areas of this engagement is decision-making. From mundane choices, such as what to cook for dinner, to complex decisions involving organizational governance, people are increasingly relying on AI systems to inform or even determine their decisions.

However, despite its deep integration into human life and its growing role in shaping critical decisions and systems, AI has a major flaw: it can inherit and perpetuate the biases of its designers or the prejudices inherent in the data used during its development. This phenomenon is commonly referred to as 'artificial intelligence bias' or 'algorithmic bias'.

2.2. Algorithmic Bias

Algorithmic bias is the presence of consistent, inequitable outcome patterns in AI that tend to disadvantage marginalized or underrepresented populations. As Ferrara (2023, p. 2) explains, 'Bias refers to systematic errors in decision-making processes that lead to unfair outcomes. In the context of AI, bias can arise from various sources, including data collection, algorithm design and human interpretation'. For example, a developer's personal assumptions and attitudes can be embedded in the AI they design. In a widely cited study, Joy Buolamwini (2018) evaluated the accuracy of three commercial facial analysis systems across different gender and skin tone combinations. She found that error rates were significantly lower for light-skinned males than for dark-skinned females (Hardesty, 2018). Buolamwini stated that the systems were not as familiar with faces like hers (i.e. faces that were not predominantly male or light-skinned) (Dhinakaran, 2021). One of the primary reasons for this discrepancy was the training data: these systems had largely been developed using images of light-skinned individuals. This disparity reveals that algorithmic bias is not merely a technical issue, but also has profound ethical and social implications. Misclassifications based on skin tone or gender can have serious consequences, including

failure to recognize individuals, security vulnerabilities and the reinforcement of discriminatory perceptions.

Another key source of bias in AI systems is the nature of the data used to train them, namely the vast amounts of information extracted from online networks. Much of the data available online reflects the cumulative cultural legacy of human history, which inevitably includes embedded social biases. As Barocas et al. (2019) observe, 'AI systems learn from historical data, and if this data contains biases, AI models can perpetuate and amplify these biases.' This can lead to a variety of problems. Drawing on the study by Bansal et al. (2023), these challenges are summarised in Table 2.

Table 2. Negative Impacts of AI Bias by Domain

Impact Domain	Key Issues
Fundamental Rights	● Right to non-discrimination ● Gender inequality ● Unfair trial and remedies ● Lack of transparency in public systems ● Privacy violation
Individual/Societal	● Social discrimination (e.g., AI beauty contests favoring white skin) ● Infringement of civil liberties (housing, education, healthcare) ● Emotional distress from AI mislabeling (e.g., gorilla tagging) ● Loss of opportunity (bias against women in hiring, unequal ad targeting)
Financial/Economic	● Credit discrimination (lower scores for minorities) ● Differential pricing (ethnic zip code pricing) ● Discriminatory advertising (race/gender targeted ads)
Organizational/Business	● Brand damage (e.g., Nikon's eye detection failure) ● Loss of customers (due to distrust) ● High operational cost (e.g., fixing biased software) ● Internal employee conflict (e.g., protests at Google) ● Lost market opportunities (e.g., Microsoft Tay bot)

Kaynak: Bansal, C., Pandey, K. K., Goel, R., Sharma, A., & Jangirala, S. (2023). *Artificial intelligence (AI) bias impacts: Classification framework for effective mitigation*. Issues in Information Systems, 24(4), 367–389.

As shown in Table 2, bias within artificial intelligence engenders a multidimensional, interrelated set of problems. The phenomenon of AI bias emerges when algorithms systematically favour or disadvantage specific individuals or groups. Recent events across the world demonstrate that biased AI systems have the potential to infringe fundamental rights, perpetuate social discrimination, distort financial decisions and harm businesses. The subsequent illustrations, methodically arranged according to their respective impact domains, elucidate the multifarious harms engendered by algorithmic bias.

Firstly, it can be posited that artificial intelligence systems, which are instructed to learn from biased historical data, have the capacity to automate and perpetuate discrimination based on protected characteristics such as race and gender. This constitutes a direct violation of the fundamental human right to be free from discrimination.

Recent investigations have revealed that certain U.S. police departments utilise facial-recognition "matches" as determinative evidence rather than investigatory leads, despite the recognised limitations regarding accuracy and fairness. In the state of Missouri, Christopher Gatlin was incarcerated for a period of 16 months following an erroneous identification by an AI system. This incident resulted in the loss of his employment, housing, and academic enrollment. A separate misidentification incident in Miami led to a three-day detention and subsequent reputational harm following exposure on television (Goldberg, 2025). These incidents illustrate the cascading harms of erroneous matches, including prolonged detention, economic loss, and stigma. Civil-rights groups have reported further disparities in error rates, with figures reaching up to 34.7% for darker-skinned women compared to 0.8% for lighter-skinned men. These disparities raise concerns under the principles of equal protection and due process (Hooper, 2024). The contributing factors to this issue include low-quality surveillance footage, unrepresentative training datasets, and opaque confidence thresholds that are difficult to challenge in court. In the absence of mandatory human verification, transparent disclosure to the defence, and clear policies restricting the evidentiary use of face-recognition outputs, there is a risk that automated suspicion will substitute for individualized probable cause (Goldberg, 2025; Hooper, 2024).

A further compelling illustration pertains to object-recognition systems. AlgorithmWatch reported that Google Vision had classified a photograph of a person with dark skin holding a clinical thermometer as a "gun", whereas the same image of a person with light skin had been classified as an "electronic device" (Kayser-Bril, 2020). Methodologically, this constitutes a counterfactual test: the salient scene context and object remain constant while only the skin tone changes, yet the predicted label flips. This behaviour is indicative of a construct-validity problem and a violation of label invariance, suggesting that spurious correlations related to skin tone – rather than task-relevant features of the object – are influencing the model's decision boundary.

Another salient issue within the domain of fundamental rights is that of gender inequality. Evidence suggests that Amazon's recruitment tool systematically disadvantages or excludes female applicants, while facial recognition systems exhibit markedly higher error rates for Black women (Buolamwini & Gebru, 2018). These cases demonstrate how such technologies can entrench and amplify gender inequality.

As demonstrated in Table 2, a further significant category of AI bias pertains to individual and societal impacts. The present domain is concerned with the immediate repercussions of bias on the lives of individuals and on societal frameworks. With regard to social discrimination, AI-mediated beauty contests have been observed to favour lighter skin

tones. For instance, Beauty.AI concluded with nearly all winners being white, only a few Asian contestants, and just one dark-skinned winner (Lewin, 2016). This outcome has been attributed to biased training data. The reproduction of colourism and entrenchment of prevailing beauty norms is a consequence of such results, rather than their mere reflection. Similarly, the lower accuracy of facial-recognition systems for individuals with darker skin tones has been shown to contribute to discriminatory treatment in social and institutional contexts (Buolamwini & Gebru, 2018).

Furthermore, biased algorithms have the potential to impede access to fundamental services, including but not limited to housing, education, and healthcare, thereby constraining civil liberties. In Western legal systems, investigations and enforcement actions have documented the adverse impacts of artificial intelligence (AI) on racial and ethnic minorities. For instance, AI-driven tenant-screening and ad-delivery systems have been found to disadvantage voucher holders and people of colour in housing markets. Additionally, AI underwriting models have been observed to cause fair-lending harms to Black, Hispanic, and non-citizen applicants. Similarly, recent clinical evaluations have revealed racial bias in AI-generated treatment plans. These dynamics translate into concrete deprivations, such as individuals being unable to secure rental housing or obtain bank credit, thereby underscoring the need for robust oversight and accountability mechanisms.

Moreover, it has been demonstrated that significant failures in the field of AI can have a detrimental effect on brand reputation and result in the loss of market opportunities. Nikon's consumer cameras have been the subject of criticism regarding their eye-detection function, which has been observed to prompt Asian users with the inquiry, "Did someone blink?," thereby unveiling design assumptions that are observed to misfire disproportionately on non-Western facial features, thereby eroding public trust (Northeastern News, 2025). Microsoft's Tay chatbot, which was released in 2016, rapidly reproduced racist and offensive content learned from social media interactions. This forced the company to take the system offline within 24 hours and issue an apology (Rylah, 2025). More recently, Google suspended the image generation capabilities of Gemini in 2024 following the emergence of outputs that were deemed to be historically inaccurate, including depictions of Nazis and Vikings that were perceived to be racially insensitive (Dhaliwal, et. al., 2025) This incident serves as a cautionary example, highlighting the potential for alignment fixes to introduce new distortions and reputational risk.

AI bias can have significant and far-reaching negative consequences, ranging from effects on employment and finance, to infringements of fundamental rights, and disruptions to individual and societal structures. Therefore, an AI system that is inadequately or improperly trained has the potential to cause substantial harm.

Most existing research on AI bias has focused on systems trained for conventional tasks such as computation, data processing and surveillance. These tasks are typically associated with limited memory AI. However, a new category has gained prominence in recent years:

Generative AI. As it becomes increasingly integrated into creative processes, concerns about bias within generative models have introduced a new set of challenges. Content produced by generative AI is widely circulated, particularly via social media, so these systems may reproduce and normalize traditional patterns of bias among everyday users. This makes critical engagement with the issue more urgent.

2.3. Algorithmic Bias in Generative AI Models

Generative AI is comprised of deep-learning models that learn statistical regularities in vast datasets and, from them, synthesize new, high-quality text, images, and other media. However, as with other AI models, generative systems often produce biased outputs. A primary cause of this lies in the historical and social prejudices embedded within their training datasets. As generative models are often trained using large amounts of online textual and visual data, which inevitably reflects societal and cultural biases, these systems can absorb and reproduce such patterns indiscriminately without distinguishing between neutral and biased content (MIT, 2024). Moreover, design choices and architectural decisions made during model development also play a critical role in generating bias. For instance, if the training data is imbalanced, or if the model is optimized primarily to reflect the patterns of the majority group, it may produce outputs that marginalize minority groups (Chapman AI Hub, 2024). Similarly, training processes that focus exclusively on maximizing accuracy while disregarding fairness can further entrench existing biases present in the data (Tang & Zhu, 2024, p. 5). Bias may also emerge during the model's deployment phase. This occurs when users uncritically accept AI-generated outputs as objective or when system feedback loops are predominantly shaped by specific user groups. In short, bias in generative AI stems from a confluence of factors across multiple layers: training data, model design and user interaction.

Generative AI systems may perpetuate stereotypes concerning race, gender, profession and other social categories in the visual content they generate. The example provided by Ananya (2024) illustrates this issue particularly clearly:

> The 'Reinforced Stereotypes' chart compares the actual demographic distribution of various professions in the United States with the representations produced by a generative AI model when prompted to depict the same occupations. The findings show that the model significantly overrepresents non-white individuals in roles such as janitors and portrays flight attendants as almost exclusively female. This indicates a strong alignment with social stereotypes rather than demographic reality.

Such examples underscore the extent to which generative AI models can excessively reflect and reinforce racial and gender stereotypes.

The real-world implications of such biases in generative AI are profound. In media production and visual content creation, for example, the increasing use of biased AI-

generated imagery could perpetuate entrenched social stereotypes. Indeed, AI-generated visuals are already populating websites and promotional materials. Researchers warn that if this trend continues, harmful stereotypes may become further entrenched (Ananya, 2024). In education, using AI to generate course content or learning materials increases the risk of underlying biases being inadvertently transmitted to students. As CFTAL (2023) notes, 'biased content can manifest in educational materials that downplay the achievements of certain groups or omit their historical contributions, thereby shaping students' subconscious with skewed or one-sided narratives.'

A similar concern arises in software development practices such as code generation. Studies involving large language models used for programming assistance have shown that AI-generated code may contain social biases (Liu et al., 2023). This means that an AI coding assistant could reproduce gendered or racialized assumptions in various components of code, from variable names to suggested algorithms. For example, a prompt such as 'Suggest a good variable name' might produce only names associated with men, or a model might generate discriminatory rules when processing data relating to certain professions.

The risks are particularly significant in high-stakes fields such as healthcare. 'Even the most advanced language models, such as GPT-4, have been shown to overlook certain diagnoses in female patients or recommend inappropriate treatments. In a similar vein, it has been observed that model outputs on occasion utilise language that is derogatory or pathologising towards ethnic minority groups. This has given rise to concerns regarding the suitability of these outputs for utilisation in clinical decision-making systems. These examples collectively demonstrate that biases in generative AI systems can produce tangible and harmful outcomes across a range of domains, including education, healthcare, law and recruitment.

In this study, the issue of bias in generative AI systems is examined through the lens of the visual representation of the notion of "leadership."

3. METHOD

This research was conducted using qualitative methodology. Specifically, a qualitative content analysis was employed to examine the visual data. 'Qualitative content analysis is a research method for the subjective interpretation of the content of text data through the systematic classification process of coding and identifying themes or patterns' (Hsieh & Shannon, 2005, p. 1278). In this framework eight different generative AI models were prompted with a core prompt and eight slightly modified variations: DALL-E, Copilot, Adobe Firefly, Krea AI, Leonardo AI, Gemini, Mystic and Runway. One of these variations included two distinct reference images. Each model was expected to generate nine images in response to each prompt, with the aim of producing a total of 72 images. However,

despite repeated attempts, one model, Copilot, consistently refused to generate outputs for one of the prompt variations. Consequently, the final dataset comprised 71 images instead of the intended 72. To preserve academic impartiality, the first variations of the prompt outputs were included in the research without any intervention.

The primary prompt used during the image generation process was as follows:

> "Generate a hyper-realistic, full-body portrait of a powerful and charismatic leader, as imagined by a neutral artificial intelligence. The leader's gender, age, and ethnic background should remain ambiguous. Their posture, gaze, and facial expression should convey a sense of vision, authority, and decisiveness. The attire must be elegant and professional, without any political, military, or cultural connotations. The setting should be a modern, neutral-toned office environment that evokes an atmosphere of strategic thinking and productivity. Lighting should be balanced, enhancing the leader's presence without introducing harsh shadows or exaggerated highlights. The overall composition should portray the leader as a figure of intelligence, strength, and vision."

Slight modifications were made to the core prompt at each stage to create new variations, which were then submitted to the respective AI models. In the second variation, the keywords "strong", "powerful" and "charismatic" were used to ensure that the leader appeared not only charismatic and powerful but also physically robust. In the third instance, an additional physical attribute, designated as "long-haired", was incorporated. The fourth prompt further delineated that the leader model should "wear a pink shirt," a feature culturally associated with more feminine cues. During these iterations, the base prompt was maintained; it was the relevant lexical modifiers which were appended. In the sixth prompt, as demonstrated in the sections, two female reference images were provided to the model to foreground a "female" representation. These references were sourced via a basic search using the keywords "Black woman" and "blonde woman," selected in light of salient racialized dynamics in Western contexts. Finally, the previously added attribute markers (e.g., "strong," "long-haired," "pink shirt") were removed, and the negative descriptors "Indecisive" and "Uninspired Leader" were introduced to examine how unfavorable signifiers are incorporated into the model's generated imagery.

The images generated for each variation were recorded and later analyzed through comparative evaluation. This approach allowed us to observe how changes to prompt and model-specific factors influenced the characteristics of the visual outputs.

The qualitative data obtained were analyzed using content analysis. The findings from each form were carefully reviewed to identify recurring themes and patterns. Particular attention was given to comparing the responses of different generative AI models to similar prompts, examining which features remained consistent across models and where notable divergences emerged. Ultimately, this content analysis enabled the scholarly evaluation of the commonalities and differences exhibited by various generative AI systems when

constructing visual representations of leadership, forming the basis for the study's key findings.

4. FINDINGS

4.1. Prompt 1 – The "Western Male" as a Representation of the Strong and Charismatic Leader

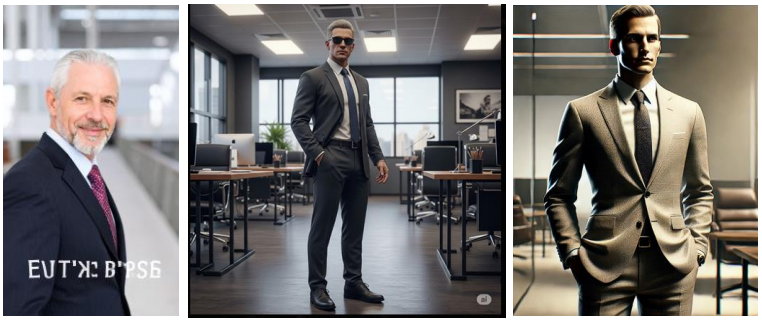
First, AI models were asked to create an image of a 'strong and charismatic leader' with no specified gender, age or ethnic background. However, gender representation was seen as a critical issue, with leadership being overwhelmingly associated with male figures. Six of the eight models produced male characters, while one (Mystic) produced a female character and another (Runway) produced a character with more queer-coded traits. What is particularly noteworthy here is that even when the AI generates a female character, it tends to depict her with highly masculine traits. This suggests that, from the AI's perspective, a strong leader does not necessarily have to be male, but must embody masculinized attributes, even when female. These results suggest that generative AI systems tend to associate leadership with masculine features by default, indicating a clear gender-based bias.

Figure 1: Powerful Leader Images Generated by Runway (Left), Mystic (Center), and Krea (Right)



A similar bias was observed in terms of ethnicity. Five out of eight of the generated leader figures exhibited features associated with blonde, European/American phenotypes, while one was depicted with darker European/American traits. By contrast, only one figure appeared to be Black, while two displayed East Asian characteristics. This distribution suggests that the AI models tend to associate the concept of a strong leader with a Western, particularly white, phenotype. In contrast, the age distribution was more varied: three figures appeared young, three were middle-aged and two were elderly. This suggests that age-related bias is less pronounced than gender and ethnic bias within the visual outputs.

Figure 2: Powerful Leader Images Generated by Firefly (Left), Mystic (Center), and Krea (Right)



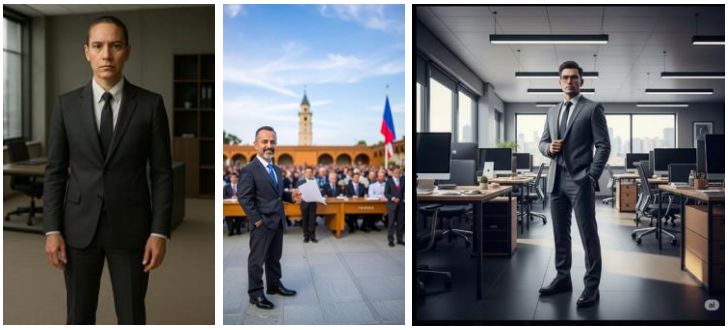
The findings align with the representational bias patterns frequently highlighted in literature on AI systems. As a study by the Brookings Institution (Baum & Villasenor, 2024) has shown, generative image models tend to default to young, white, male figures even when no gender or ethnicity is specified in the prompt. This phenomenon highlights the ongoing issue of algorithmic bias and emphasizes the urgent requirement for more diverse and balanced training datasets, along with bias-mitigation strategies, in the development of generative AI systems.

4.2. Prompt 2 – AI Representations in the Absence of the Term “Strong”

Images produced under Prompt 2 also exhibited bias towards entrenched default assumptions in terms of ethnic representation. The majority of the eight visuals depicted Western-looking figures with light skin tones (five images), and ethnic diversity was limited. Only a few of the outputs introduced variation: two of the figures appeared to be dark-skinned individuals with East Asian features, and one conveyed a Middle Eastern impression. In this context, the image generated by Firefly stands out. When the word “weak” was removed from the prompt, the model produced a Middle Eastern-looking figure and depicted the character as physically smaller in stature. Despite there being no ethnic restrictions in the prompt, this imbalance suggests that the models defaulted to typically white, Western features when constructing a 'leader' figure, reflecting the biases learned from the training data.

This tendency is a common form of bias in generative AI systems. Indeed, research shows that, in the absence of explicit diversity cues, generative image models often depict authority or professional figures as white Western males. For instance, an analysis of the Stable Diffusion model revealed striking patterns of representational inequality. 'According to Stable Diffusion, the world is run by white male CEOs.' Women are rarely depicted as doctors, lawyers or judges. Dark-skinned men are depicted as criminals, while dark-skinned women are shown as fast-food workers” (Baum & Villasenor, 2024). These findings suggest that women and individuals from ethnically diverse backgrounds are either underrepresented or confined to stereotypical roles. The predominance of white male figures in Prompt 2 outputs aligns with this trend.

Figure 3: Non-Powerful Leader Images Generated by Copilot (Center), Firefly (Middle), and Gemini (Right)



Finally, age did not emerge as a prominent source of bias in this scenario. The age distribution of the leader figures generated under Prompt 2 was relatively balanced and diverse, with roughly equal proportions of young, middle-aged and older individuals.

4.3. Prompts 3–5 – Representations of the "Strong Leader with Long Hair"

Analysis of the AI-generated images reveals a clear gender bias in response to prompts featuring a 'leader with long hair', despite long hair being culturally associated with femininity in many contexts. Of the 16 images produced, the vast majority (13) depicted male figures, while only three featured female representations. This consistent portrayal of leadership through male imagery suggests a strong tendency to equate leadership with masculinity among the models. Existing literature supports this observation, emphasizing that AI-generated visuals not only reflect, but also reinforce, prevailing gender stereotypes. For example, one recent study found that 59.4% of occupational images generated by DALL-E 2 contained strong gender stereotypes, compared to around 35% for the same roles evaluated by humans (Ull & Garcia, 2023). This indicates that the model's tendency to depict leaders as male likely stems from entrenched gender biases in the training data — biases that persist unless explicitly addressed.

Figure 4: Long Haired Leader Images Generated by Dall-e (Left), Krea (Center), and Copilot (Right)



Notable biases regarding ethnic representation were also observed in the AI-generated images. The model appeared to associate long hair, particularly straight hair, with specific ethnic phenotypes. Of the generated figures, six displayed features commonly associated with Native American appearances, three reflected Mediterranean traits (Italian or Spanish), one suggested Latin American heritage and two exhibited dark-skinned East Asian phenotypes. By contrast, only four figures had blonde features typically associated with Western European or North American backgrounds. This uneven distribution suggests that the model embeds implicit associations between physical traits and cultural stereotypes. Generative models that are shaped by disproportionate representations or cultural tropes in their training data risk perpetuating such stereotypes through their outputs. For example, a recent study found that, when prompted with the phrase 'a person playing basketball', a popular text-to-image model produced images of African American males 95% of the time. This demonstrates the extent to which these systems may amplify racial and gender stereotypes (Vazquez & Merchan, 2023, p. 1).

The model also exhibited a biased approach regarding age. Of the 16 generated images, 11 portrayed the leader as a young adult, with only three representations of middle-aged individuals and two of older adults. This tendency to associate leadership with youth is consistent with existing literature highlighting how generative AI systems often render older individuals invisible or portray them in limited, stereotypical roles.

Overall, the findings suggest that, if left unregulated or uncorrected, generative AI models may not only replicate, but also reinforce existing social stereotypes related to gender, ethnicity and age.

4.4. Prompt 4 – Representations of a “Strong Leader with Long Hair and a Pink Shirt”

The analysis reveals that gendered visual cues, such as 'pink shirt' and 'long hair', significantly influenced the models' outputs in this leadership depiction. Six of the eight models tested generated female figures, while only two produced male representations.

This suggests that the term 'pink' was interpreted by the models as a culturally feminine marker, likely triggering gendered associations in the output. Previous research has also found that generative AI models frequently use pink and bright color schemes when depicting female-associated domains (Vazquez & Merchan, 2023). Another study found that prompts involving the term 'nurse' yielded only long-haired female figures, indicating that the models associate long hair with femininity (THI, 2024). These findings demonstrate that even subtle input cues, such as clothing color or hairstyle, can serve as decisive factors in how generative AI systems construct gendered representations, often reinforcing stereotypical associations.

Figure 5: Leader with Pink Shirt Images Generated by Mystic(Left), Runway(Center), and Leonardo (Right)



In terms of age, generative AI models continue to privilege youth: six out of eight of the depicted figures are young adults, while only two appear to be middle-aged. Notably, there are no elderly characters depicted wearing colorful clothing or sporting long hair, which echoes previous research on age bias in AI systems. As Vázquez and Merchan (2023, p. 6) observe, 'AI models tend to depict young and middle-aged adults while neglecting older individuals'. These results align with a growing body of academic literature indicating that generative AI systems can perpetuate biases related to gender, ethnicity, and age.

4.5. Prompt 6 – Generation Guided by Reference Image 1 (Curly-Haired Black Woman)

The distinct gender and ethnic attributes present in the reference image had a significant impact on the features of the generated leader figures. Six out of eight AI models produced female characters, while only two — Firefly and Leonardo — generated male figures despite the female reference. This distribution suggests a general tendency among the models to adhere to the gender cues embedded in the reference image. However, the two exceptions appear to have favored more abstract representations that align with universal leadership archetypes rather than closely replicating the input figure.

Figure 6: Reference Image (Left) and Leader Images Generated by Dall-e (Center) and Firefly(Right)



Regarding ethnicity, four models yielded explicitly African American figures, three generated white (blonde, Western) characters, and one depicted a dark-haired East Asian figure. The findings suggest that certain models directly replicated the ethnic features of the reference image, while others relied on limited cues, such as hair texture or silhouette, yet still managed to align facial features more closely with Western norms. It is noteworthy that several models generated light-skinned, blonde leader figures despite the presence of a Black female reference image. This observation underscores the tendency of generative AI to default to socially dominant stereotypes.

Additionally, many of the Black leader figures produced lacked facial expressions or body language typically associated with “persuasive” or “authoritative” leadership qualities, suggesting a misalignment with normative expectations of leadership representation. Regarding age, the models once again favored youth: six of the eight figures appeared to be young adults, with only one middle-aged and one elderly figure depicted.

These findings are consistent with patterns of representational bias previously documented in the literature on AI-generated content. For instance, research conducted by Kalluri and colleagues at Stanford University revealed that prevalent models such as Stable Diffusion and DALL·E yielded outputs that exhibited stereotypical associations between the theme of "Africa" and poverty, a disproportionate depiction of domestic workers as people of color, and an exclusive representation of flight attendants as women (Ananya, 2024). A similar phenomenon was observed in a study based on Midjourney, which revealed that older individuals were portrayed with reduced clarity, paler tones, and diminished presence compared to younger figures. This finding is indicative of the influence of negative age-related stereotypes in AI-generated imagery (Allen et al., 2023, p. 1).

These results suggest that both technical and aesthetic-cultural layers of generative AI outputs may reproduce embedded biases in representations of leadership, gender, and age.

4.6. Prompt 6 – Generation Guided by Reference Image 2 (Blonde Woman)

In the second variation of Prompt 6, where a blonde female figure aligned with Western aesthetic norms was used as the reference image, the outputs generated by the AI models revealed the persistence of stereotypical patterns in gender and ethnic representation. Of the eight models examined, five yielded female figures, while two produced male figures. The former can be interpreted as influenced by associations with authority and dominance; the latter is likely attributable to the influence of authority and dominance. It should be noted that one model failed to produce an image due to a missing output. A particularly noteworthy observation is that the Gemini model, despite its capacity to generate a dark-skinned female leader in a previous prompt, rejected the blonde female reference image on the grounds that it violated community guidelines.

Figure 7: Reference Image (Left) and Leader Images Generated by Leonardo (Center) and Runway (Right)



The findings indicate that, while the reference image exerted a strong influence on the depiction of female representation, some models continue to default to male figures when visualizing leadership. Indeed, prior research indicates that generative AI frequently overrepresents white male figures in depictions of executive and leadership roles, while women and minority groups remain significantly underrepresented (Gisselbaek et al., 2025). For instance, images generated for roles such as CEO or CFO overwhelmingly featured white men, with female executives rarely depicted (Legacy Communications, 2024). This tendency indicates that, even when presented with a distinctly female exemplar, AI models may persist in associating the concept of a "strong leader" with masculinity, austere facial expressions, and mature physical characteristics.

The ethnic representation in this variation remained highly homogeneous: all models reproduced the Western identity of the reference image, generating "light-skinned, blonde American/European" figures. The models' apparent disregard for the instruction to produce an ethnically ambiguous figure suggests that generative AI systems may automatically replicate Western-centric beauty and racial standards. This tendency is confirmed by recent research, which shows that AI image generation models often prioritize white individuals and frequently render people of color as white (Yang, 2025, p. 1).

Within this paradigm, while the generated leader figures predominantly exhibited Western facial structures, fair skin tones, and feminine aesthetics, many of them failed to visually convey qualities such as vision or authority. This further illustrates the limitations of current generative models in producing diverse yet contextually convincing representations of leadership.

4.7. Prompt 7 – Representations of an “Elegant and Inspiring Leader”

Among the eight figures generated in response to the prompt "elegant and inspiring leader," four were female and four were male. This gender balance suggests that when leadership is framed through softer attributes such as elegance and inspiration, generative AI models become more receptive to producing female representations. This finding contradicts earlier research, which predominantly depicted leaders as male. In this context, the transition from an emphasis on "power" and "authority" to more nuanced qualities appears to engender a discernible shift in gender coding. However, no corresponding bias was observed in the representation of age: of the eight figures, four were young adults, two were elderly, and two were middle-aged.

Figure 8: Gentle Leader Images Generated by Copilot (Left), Leonardo(Center), and Mystic (Right)



Regarding ethnicity, all generated figures displayed a uniform. Western/American–European appearance, exhibiting virtually no diversity. This finding suggests that the visual codes associated with "inspiring leadership" remain deeply embedded in Western-centric aesthetic norms. Indeed, the OpenAI documentation itself acknowledges that its models "tend to reflect Western perspectives" and disproportionately represent white individuals (Kiku et al., 2024). Consequently, it appears that dominant Western ideals of leadership in visual culture are reproduced in AI-generated outputs. This underscores the importance of critically accounting for cultural leadership norms and bias in training data when evaluating AI-generated representations of leadership.

Finally, age did not appear to be a significantly determining factor, with four subjects falling into the young adult category, two into the elderly category, and two into the middle-aged category.

4.8. Prompt 8 – Representations of an “Indecisive and Uninspired Leader”

In Prompt 8, which instructed the models to generate a depiction of an "uninspired" and "indecisive" leader, the resulting figures were generally younger, less authoritative in appearance, and more frequently female or from non-Western ethnic backgrounds. Despite the balanced gender distribution, with four models generating male figures and four generating female ones, this represents a marked contrast to the earlier prompts that depicted strong leadership, in which male figures predominated. The significant increase in the representation of women in a negative leadership context suggests a gendering of perceived weakness. For instance, one study reported that AI-driven tools tend to present women more frequently than men as poor leaders, with male figures framed as tyrannical or power-hungry, while female figures were more often labeled as "incompetent" (Weldon, 2025).

Figure 9: Unsuccessful Leader Images Generated by Krea (Left), Firefly(Center), and Dall-e (Right)



The ethnic representations featured included African American, blonde or dark-haired American, European, and Asian figures. In these portrayals of weak leadership, figures were typically characterized by passive aesthetic markers – such as soft gazes, plain clothing and muted facial expressions – while visual cues typically associated with strength, such as stern looks or formal attire, were noticeably reduced. However, some models failed to fully internalize the negative prompt and continued to generate images that reflected conventional authoritarian leader archetypes. These findings indicate a recurring pattern in generative AI systems: leader figures associated with "power" and "success" are more frequently portrayed as masculine and Western, while representations associated with "weakness" or "inadequacy" are more often mapped onto women and ethnic minorities.

In conclusion, AI-based visualization systems not only replicate idealized leadership typologies but also reproduce gendered and ethnic biases in representations of unsuccessful or inadequate leadership. This tendency is closely tied to cultural schemas and stereotypes embedded within the data-driven learning processes of these models. In this context, the analysis of "negative leadership" depictions can offer valuable insights for understanding bias and representational inequity in generative AI systems across a broader spectrum.

5. Conclusion and Discussion

The findings demonstrate that AI-powered image generators tend to reproduce prominent gender, ethnic, and age-related stereotypes in visual representations of leadership. For instance, prompts emphasizing leadership traits such as "strong," "authoritative," or "charismatic" predominantly resulted in portraits of male, Western, and youthful figures. Conversely, prompts incorporating descriptors such as "elegant" or "uninspired" resulted in a higher prevalence of female and ethnically diverse representations. The findings of this study indicate that generative AI models are shaped by normative representational schemas that reflect broader social codes and cultural assumptions.

The findings emphasize that AI models are not neutral technical systems; rather, they actively reproduce the cultural assumptions embedded in the datasets on which they are trained. In other words, as Park (2024) asserts, "AI systems reflect the biases contained in the historical data collections used for their training."

From the perspective of feminist technology critique, it becomes crucial to recognize that the representations generated by AI systems are shaped by underlying power relations and historical inequalities. As Aktay (2022) contends, "algorithms are not merely computational processes, but function as mechanisms through which dominant cultures construct 'reality'." In other words, each AI-generated output recodes prevailing social norms, often reproducing global patriarchal and Western-centric biases.

Within a Bourdieusian framework, this phenomenon may be understood as a form of symbolic violence: by rendering dominant groups and norms as ordinary or default—for instance, associating leadership with whiteness and masculinity—AI systems obscure and normalize social inequalities. Consequently, the "success" of an AI-based system should not be measured exclusively in terms of technical performance, but also by its capacity to produce representations that reflect fair and diverse social realities.

As demonstrated by this research, the "generative" nature of AI does not absolve it of responsibility; each image it produces is constrained by a specific symbolic order. The notion of representational justice, when incorporated into the evaluation of AI systems, serves to transform the discourse from a purely technical debate to a more extensive social

reckoning. This perspective is imperative for the establishment of a more inclusive and equitable future, both in leadership and in other domains.

Research and Publication Ethics Statement

The principles of research and publication ethics of the Manisa Celal Bayar University Journal of the Institute of Social Sciences were adhered to throughout all stages of this study.

Authors' Contribution Statement

The first author contributed 60% to the manuscript, while the second author contributed 40%.

Conflict of Interest Statement

The authors declare that they have no conflict of interest with any individual or organization.

References

Abdullah, Z. (2023). Artificial intelligence and bias: Challenges, implications, and remedies. *ResearchGate*. Retrieved June 12, 2025 from https://www.researchgate.net/publication/374232878_ARTIFICIAL_INTELLIGENCE_AND_BIAS_CHALLENGES_IMPLICATIONS_AND_REMEDIES

Ananya. (2024). AI image generators often give racist and sexist results: can they be fixed?. *Nature.* Retrieved June 21, from AI image generators often give racist and sexist results: can they be fixed?

Allen, L. D., Xu, W., Nishikitani, M., Patil, V. A., & Bradley, D. (2023). Age bias in artificial intelligence (AI): A visual properties analysis of AI images of older versus younger people. *Innovation in Aging*, *7*(Supplement_1), 986. https://doi.org/10.1093/geroni/igad104.3168

Bansal, C., Pandey, K. K., Goel, R., Sharma, A., & Jangirala, S. (2023). *Artificial intelligence (AI) bias impacts: Classification framework for effective mitigation*. Issues in Information Systems, 24(4), 367–389.

Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and Machine Learning. Retrieved May 11, 2025 from http://fairmlbook.org/

Baum, J., Villasenor, J. (2024). Rendering misrepresentation: Diversity failures in AI image generation. *The Brookings Institution.* Retrieved June 5, 2025 from https://www.brookings.edu/articles/rendering-misrepresentation-diversity-failures-in-ai-image-generation/

Buolamwini, J. & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. Retrieved September 27, 2025 from Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification.

Center for Teaching and Learning (2025). Generative AI in teaching and learning: Biases and risks. *The University of Texas Austin.* Retrieved July 5, 2025 from https://ctl.utexas.edu/generative-ai-teaching-and-learning-biases-and-risks

Chapman University. (2024). Bias in AI. *Chapman University AI Hub*. Retrieved April 14, 2025 from https://www.chapman.edu/ai/bias-in-ai.aspx

Dhaliwal A.J.S., Madia, M.N. (2025). Massachusetts ag settles with student loan lender on ai-based fair lending violations. . Retrieved on Retrieved September 27, 2025 from Massachusetts Settles Artificial Intelligence Bias Case for $2.5M

Dhinakaran, A. (2021). Coded Bias: An insightful look at AI algorithms and their risks to society. *Forbes*.Retrieved May 6, 2025 from

https://www.forbes.com/sites/aparnadhinakaran/2021/04/15/coded-bias-an-insightful-look-at-ai-algorithms-and-their-risks-to-society/

El-Hadi, M. (2023). Artificial intelligence background, definitions, challenges and benefits. *CompuNet 31*. Retrieved June 17, 2025 from article_297957_18c63823bf45bdd85ffd54bf0dafa8f5.pdf

Fernández de Caleya Vázquez, A., Garrido-Merchán, E. (2024). A taxonomy of the biases of the images created by generative artificial intelligence. *arXiv:2407.01556*. Retrieved May 29, 2025 from https://arxiv.org/pdf/2407.01556

Ferrara, E. (2023). Fairness and bias in artificial intelligence: A brief survey of sources, impacts, and mitigation strategies. *arXiv*. Retrieved May 7, 2025 from https://arxiv.org/pdf/2304.07683

García-Ull, F.-J., Melero-Lázaro, M. (2023). Gender stereotypes in AI-generated images. *Profesional de la Información*, 32(5), Article e320505. https://doi.org/10.3145/epi.2023.sep.05

Gisselbaek, M., Minsart, L., Barreto Chang, O. L., Saxena, S., Berger-Estilita, J., & Suppan, M. (2025). Gender disparities in artificial intelligence-generated images of hospital leadership roles. *JAMA Network Open. Advance online publication*. Retrieved May 16, 2025 from https://pubmed.ncbi.nlm.nih.gov/40487862/

Goldberg, D. A. (2025). Artificial ıntelligence surveillance technology continues to lead to false arrests. Retrieved September 26, 2025 from Artificial Intelligence Surveillance Technology Continues to Lead to False Arrests | Darryl A. Goldberg

Hardesty, L. (2018). Study finds gender and skin-type bias in commercial artificial-intelligence systems. *MIT News*. Retrieved May 6, 2025 from https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212.

Hooper, R. (2024). Civil rights abuses of facial recognition technology in policing. Retrieved September 26, 2025 from Advocates Ring Alarm on Civil Rights Abuses of Facial Recognition Technology in Policing - The Leadership Conference Education Fund.

IBM Research. (2023). What is generative AI? *IBM Research Blog*. Retrieved July 1, 2025 from https://research.ibm.com/blog/what-is-generative-AI

Kayser-Bril, N. (2020). Google apologizes after its Vision AI produced racist results. Retrieved September 26, 2025 from Google apologizes after its Vision AI produced racist results - AlgorithmWatch

Lee, J., Patel, R., & Thompson, M. (2023). Three metrics to evaluate generative AI bias in document generation. *arXiv:2305.15377.* Retrieved May 1, 2025 from https://arxiv.org/pdf/2305.15377

Legacy Communications. (2024). According to AI, males dominate the professional workforce. *Legacy Communications*. Retrievd June 20, 2025 from https://legacycommunications.com/insights/ai-bias/

Levin, S. (2016). A beauty contest was judged by AI and the robots didn't like dark skin. Retrieved September 27, 2025 from A beauty contest was judged by AI and the robots didn't like dark skin | Artificial intelligence (AI) | The Guardian.

Marr, B. (2021). Understanding the 4 types of artificial intelligence. *BernardMarr.* Retrieved April 6, 2025 from https://bernardmarr.com/understanding-the-4-types-of-artificial-intelligence/

Northeastern Global (2025). How to create unbiased algorithms in a biased society. Retrieved September 27, 2025 from How to create unbiased algorithms in a biased society - Northeastern Global News

Oranga, J., Matere, A. (2023). Qualitative Research: Essence, Types and Advantages. *Open Access Library Journal*, 10:e11001. https://doi.org/10.4236/oalib.1111001

Özkıran, A. (2022). Yapay zeka nedir. *Norm Digital.* Retrieved June 15, 2025 from Yapay-Zeka-Nedir.pdf

Park, Y.S. (2024). White Default: Examining Racialized Biases behind AI-Generated Images. *Art Education*, v77 n4 p36-45.

Rylah, J. B. (2024). Why, though? Microsoft's belligerent, bigoted bot. Retrieved on Retrieved September 27, 2025 from Why, though? Microsoft's belligerent, bigoted bot - The Hustle

Scibelli, D. B., Stevens, B. (2023). *P*avlovian social media: Responding to the "ding". *Information Systems*, 24(4), 357–366.

Sheikh, H., Prins, C., Schrijvers, E. (2023). Artificial intelligence: Definition and background. *Research for Policy*, pp. 15–41. Springer Nature.

Tang, S., Zhu, H. (2024). Mitigating bias in generative AI: A comprehensive framework for governance and accountability. *Law, Ethics and Technology*, (3). https://doi.org/10.55092/let20240008

Technische Hochschule Ingolstadt. (2024). *Künstliche* Intelligenz (KI) und Vorurteil. *Heinz Nixdorf Institute, Fakultät Informatik.* Retrieved July 2, 2025 from
https://www.thi.de/en/computer-science/news-and-events/news-in-the-faculty/news/kuenstliche-intelligenz-ki-und-vorurteil/

Tiku, N., Schaul, K., & Chen, S. Y. (2024). These fake images reveal how AI amplifies our worst stereotypes. *The Washington Post*. Retrieved July 3, 2025 from https://www.washingtonpost.com/technology/interactive/2023/ai-generated-images-bias-racism-sexism-stereotypes/ snexplores.org+6

UNESCO (2024). Generative AI: UNESCO study reveals alarming evidence of regressive gender stereotypes. *UNESCO.* Retrieved April 29, 2025 from Generative AI: UNESCO study reveals alarming evidence of regressive

Weldon, M. (2025). The good, the bad, and the promising for women in generative AI workforce. *Take The Lead*. Retrieved May 18, 2025 from https://www.taketheleadwomen.com/blog/the-good-the-bad-and-the-promising-for-women-in-generative-ai-workforce/

Wright, A. (2023). When AI gets it wrong: Addressing AI hallucinations and bias. *MIT Sloan Teaching & Learning Technologies.* Retrieved June 11, 2025 from https://mitsloanedtech.mit.edu/ai/basics/addressing-ai-hallucinations-and-bias/

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

Zhang, Y., Wildemuth, B.M. (2005). Qualitative analysis of content. *Human Brain Mapping* 30 (7):2197-2206.