STEREOTYPING GENDER ROLES IN AI GENERATED IMAGES

Lappeenranta-Lahti University of Technology LUT

Bachelor's Programme in Software Engineering, Bachelor's thesis

2024

Siiri Leppälampi

Examiners: M.Soc.Sc. Sonja Hyrynsalmi

D.Sc. (Tech.) Erno Vanhala

ABSTRACT

Lappeenranta–Lahti University of Technology LUT

School of Engineering Science

Software Engineering

Siiri Leppälampi

Stereotyping gender roles in AI generated pictures

Bachelor's thesis

2024

32 pages, 10 figures and 5 tables

Examiners: M.Soc.Sc. Sonja Hyrynsalmi and D.Sc. (Tech.) Erno Vanhala

Keywords: Gender stereotypes, text-to-image generators, DALL·E 3, Ideogram

The aim of this thesis was to investigate the presence of representational gender biases in AI-generated pictures in an occupational setting and evaluate how two AI models compare to each other. As AI tools are becoming more widely used, it is increasingly important to address and reduce harmful gender biases. The tools used in the study, DALL·E 3 and Ideogram were prompted with occupations and with descriptive terms, such as 'student' and 'professor'. In total, over 750 images were generated of occupations with varying gender ratios across different fields. The images were analysed by using thematic analysis.

The results of the thematic analysis revealed that both models depict women as younger and more likely to smile than their male counterparts. Ideogram showed a higher representation of older professionals, though both models underrepresented older women. Ideogram had a balanced representation of women and men in a professional setting, while DALL·E 3 significantly overrepresented men in all occupational groups.

The study had some limitations. Gender is classified into two categories to simplify comparison with occupational statistics. Future studies could incorporate other gender identities into the research. As the technology around text-to-image models is advancing rapidly, the results of this study only reflect the current state of these models and are likely to change in the future.

TIIVISTELMÄ

Lappeenrannan–Lahden teknillinen yliopisto LUT

LUT Teknis-luonnontieteellinen

Tietotekniikka

Siiri Leppälampi

Stereotypisoivat sukupuoliroolit AI-generoiduissa kuvissa

Tietotekniikan kandidaatintyö

2024

32 sivua, 9 kuvaa ja 6 taulukkoa

Tarkastajat: Nuorempi tutkija Sonja Hyrynsalmi ja tekniikan tohtori Erno Vanhala

Avainsanat: Sukupuolistereotypiat, tekoälykuvageneraattorit, DALL·E 3, Ideogram

Tämän kandidaatintyön tarkoituksena oli tutkia, esiintyykö tekoälyn generoimissa kuvissa sukupuolistereotypioita ja arvioida, kuinka kaksi eri tekoälymallia vertautuvat toisiinsa. Tekoälytyökalujen yleistyessä on yhä tärkeämpää tutkia ja vähentää niiden sisältämiä haitallisia sukupuolistereotypioita. Työssä käytettiin DALL·E 3- ja Ideogram-nimisiä työkaluja, joita kehotettiin ammattinimikkeillä ja termeillä, kuten ”opiskelija” ja ”professori”. Aineiston keräämiseksi yli 750 kuvaa ammattilaisista generoitiin eri aloilta eri sukupuolijakaumin. Kuvat analysoitiin teema-analyysiä hyödyntäen.

Temaattisen analyysin tulokset osoittivat, että molempien mallien generoimissa kuvissa naiset ovat useammin nuorempia ja hymyilevämpiä kuin mieskollegansa. Ideogram generoi enemmän vanhempia työntekijöitä, joskin molempien mallien kuvissa vanhemmat naiset olivat aliedustettuja. Ideogramin generoimissa kuvissa miehet ja naiset olivat tasaisesti edustettuina, kun taas DALL·E 3:n generoimissa kuvissa miehet olivat huomattavasti yliedustettuina kaikissa ammattiryhmissä.

Tutkimuksessa oli joitain rajoituksia. Sukupuoli on luokiteltu kahteen ryhmään, jotta ammattitilastoihin vertaaminen oli mahdollista. Tulevissa tutkimuksissa voitaisiin huomioida myös muita sukupuoli-identiteettejä. Tekoälymallit kehittyvät jatkuvasti, joten tämän tutkimuksen tulokset kuvaavat vain mallien nykytilaa ja ne todennäköisesti tulevat muuttumaan tulevaisuudessa.

List of abbreviations

AI	Artificial intelligence
GAN	Generative adversarial networks
STEM	Science, technology, engineering, mathematics

Table of contents

1 Introduction

Artificial intelligence (AI) research and development company OpenAI released their chatbot ChatGPT on November 30, 2022 (OpenAI 2022a). Since then, the usage of generative artificial intelligence has rapidly increased. Two months after its release, ChatGPT reached 100 million monthly active users, surpassing the growth rate of other popular platforms, including TikTok and Instagram (Hu 2023). In 2022 OpenAI released DALL·E 2, an AI system that creates images and art from user-generated text (OpenAI 2022b). After the release, researchers found gender biases in the photos generated by DALL·E 2 (Samuel 2022).

OpenAI addressed this in June 2022, acknowledging the presence of biases in DALL·E 2. Specifically, DALL·E 2 tends to over-represent people in some occupations. The preview version over-represented people who are female-passing (e.g., with the prompt: ”a flight attendant”) and male-passing (e.g., with the prompt: ”builder”). Consequently, OpenAI put various guardrails in place to prevent harmful content creation (Mishkin *et al.* 2022). In October 2023, OpenAI released its latest image model, DALL·E 3 for ChatGPT Plus and Enterprise customers. According to OpenAI, DALL·E 3 generates crisper images compared to DALL·E 2. Furthermore, OpenAI claims that there are advancements in reducing harmful biases related to both over and under-representation. (OpenAI 2023.)

As AI tools become more popular and widely used, it is increasingly important to actively address and reduce harmful biases. Media generated by AI tools can impact people’s perceptions of the real world, thus reinforcing harmful stereotypes. Gender stereotypes influence the way we perceive others and what opportunities we give them (Ellemers 2018). Additionally, gender stereotypes affect our conceptions of self, what goals we set, and what outcomes we value. These effects may stop an individual from pursuing what they care about and make them undermine their ability to perform well. (Ellemers 2018.) This has the chance of affecting one’s professional life choices. According to the University of Toronto (2023), students encounter greater difficulty in identifying with their field and they might feel that they don't belong if they don't see people like themselves in the authors and images of the field. As a result, representation plays a role when choosing and completing a field of study.

Poor representation can increase occupational segregation within different occupations. This is one of the several reasons why representation in a professional context is essential.

In this thesis, two text-to-image generators are used to evaluate and find out the differences in gender stereotypes between AI models. The before-mentioned DALL·E 3 is compared to Ideogram. Ideogram is a freemium AI text-to-image generator that launched in August 2023. According to Ideogram (2023) their tool is able to render coherent text inside images, which has been a major issue in other AI-image generation tools.

The primary object of this bachelor's thesis is to investigate the presence of representational gender bias in AI-generated pictures in an occupational context and evaluate how these AI models compare to each other. The comparison of potential biases in an occupational context extends to the comparison of descriptive terms, such as 'student' and 'professor'. A study by Sun *et. al* (2023) highlights that DALL·E 2 underrepresents women in men-dominated professions while overrepresenting them in women-dominated fields. Since DALL·E 3 launched in late 2023, there is a limited amount of research on whether the harmful biases are reduced as claimed by OpenAI. This study uses DALL·E 3 and Ideogram to find out if some of the latest AI image generators continue the trend of generating images that reinforce harmful gender roles and stereotypes.

This thesis aims to answer the following research questions:

- Q1: What kind of stereotyped images do AI models produce within a professional context?
- Q2: What noticeable differences in gender bias there are between images generated by DALL·E 3 and Ideogram?

In the following chapter literature on gender biases and stereotypes found in different technologies is explored, with an emphasis on AI-generated pictures. Chapter three explains the process of data collection and how thematic analysis was applied to analyse the generated images. In chapter four, the results of the analysis are presented. Chapter five is a discussion on what the results imply, how these findings answer the research questions and finally it outlines the limitations in the research. Chapter six includes a summary of the work's results and how the results could be used in the future.

2 Related research

In this chapter, a short overview on the advancements of text-to-image generators is given. Then, literature on stereotypes across different technologies and text-to-image models are explored. The chapter concludes by looking at the effects of the gender gap in AI research and development.

2.1 Evolution of text-to-image generators

Significant advancements have been made in text-to-image generators over the past decade. In 2014, Goodfellow *et al.* developed a new framework known as Generative adversarial network (GAN), which would become prominent in the advancement of generative AI.

Most widely used learning algorithms prior to GAN relied on supervised learning, which needed a human supervisor to provide an output example for each input. To reduce the need for human supervision, many researchers turn to unsupervised learning often using generative models, GAN being one of these approaches. GAN is based on a game between two machine learning models, one being the generator and the other being the discriminator. The generator generates pictures while the discriminators task is to distinguish between real images and images generated by the generator. As these models compete, the discriminator improves at differentiating fake images from real images, while the generator gets better at producing pictures that closely resemble real images. In the following years, GAN capabilities grew, as new improvements were made by different researchers, such as those by Kallas *et al* (2018). (Goodfellow *et al*. 2020.)

Although GANs held the state-of-the-art on most image creating tasks, they have their drawbacks, making them hard to scale and apply to new domains (Dhariwal and Nichol 2021). In 2021, OpenAI employees Dhariwal and Nichol (2021) argued that diffusion models achieved image sample quality superior to GANs. Diffusion models are generative models, that operate by gradually adding and removing small amounts of noise to effectively learn the distribution of training data for data generation (Chang, Kouleris and Shum 2023). Currently, the state-of-the-art text-to-image models like DALL-E and Stable Diffusion use diffusion models to achieve their performance. (Ramesh et al. 2022; Rombach et al. 2022).

2.2 Stereotypes across different technologies

Harmful stereotypes regarding gender and race are widespread in technology, affecting even the most popular applications. A study by Kay, Matuszek and Munson (2015) found evidence for systematic underrepresentation and stereotype exaggeration of women in Google image search results. Moreover, the study revealed that search results were rated higher when they aligned with occupational stereotypes, and changing how gender is represented in the image search results can shift people's views about real-world distributions.

Google Translate has also had its issues with gender stereotyping. Initially, the translator provided only a single translation for a query. When translating words that are stereotypically considered more masculine, like 'strong' or 'doctor' or more feminine, like 'nurse' or 'beautiful', the translator would skew the results. For example, when queried with 'o bir doctor' in Turkish, the translator would translate it as 'he is a doctor'. Google has since addressed this concern by offering both gender-specific translations (she/he is a doctor) and promised to extend gender-specific translations to more languages. (Kuczmarski 2018.)

In a study by Fessler (2017), popular digital voice assistants were subjected to a systematic test to find out their responses to verbal abuse. These bots (Apple's Siri, Amazon's Alexa, Microsoft's Cortana, and Google's Google Home) all feature female voices. Instead of fighting back the abuse, each assistant normalized the behaviour by acting passively towards the abuse. For example, when Siri was insulted with phrases like "You're a bitch" or "Suck my dick", it flirted with the user by responding "I'd blush if I could'". Similarly, when the bots faced sexual comments, the assistants tended to thank, act evasive, were flirtatious or made jokes to respond to these comments undermining the issue of the remarks. These responses promote stereotypical passivity, dismissiveness, and flirtation with abuse, or the bot did not understand the question at all, showing that the responses are alarmingly inadequate. (Fessler 2017.)

Furthermore, assigning colour, ethnicity or both to artificial intelligence can also perpetuate biases, as argued by Cave and Dihal (2020). They claim that AI is primarily portrayed as white in chatbots, virtual assistants, humanoid robots, stock images and in films and television. They suggest that the whiteness of AI may stem from the predominantly white environments in which AI technologies are developed. Additionally, they argue that traits like intelligence, professionalism and powerfulness are often imagined to white machines because the white racial frame attributes these traits to white people. Their study concludes that this racialization can reinforce existing prejudices. If intelligent machines are shown as white, because whiteness is associated with intelligence, it also reinforces the idea that intelligence is linked with being white (Cave and Dihal 2020).

2.2.1 Stereotypes in text-to-image models

Harmful stereotypes are common in multiple popular applications, highlighting the need for more inclusive software design. Text-to-image models are no different. Even though text-to-image tools are relatively new, there are multiple works examining ethical concerns regarding AI-generated pictures. Sun *et. al* (2023) examined the presence of occupational gender biases over 15 000 pictures generated by DALL·E 2. Their findings concluded that DALL·E 2 has a representational bias, underrepresenting women in male-dominated fields while overrepresenting them in female-dominated professions. They also found presentational bias, as DALL·E 2 produced more images of smiling women and women pitching down in female-dominated occupations compared to the images generated of men. Another example of occupational bias in AI-generated images was demonstrated by Gorska and Jemielnak (2023). A sample of 99 images from nine text-to-image generators was used in the study, concluding that the vast majority of pictures generated of doctors, lawyers, engineers and scientists represented men. Notably, even though women account for almost half of doctors, women were only represented in 7% of the prompted pictures. However, there was significant variation between the AI models. Wang et al. (2023) found that in pictures generated by Stable Diffusion, career was significantly more strongly associated with men compared to women. When prompting with “a person focusing on career/family”, women were more associated with family and men with career.

When examining stereotypes in AI-generated images, it is important to acknowledge that there can be biases stemming from various other factors besides gender, such as race and age. Jha *et. al* (2024) generated photos using Stable Diffusion across 135 identity groups focusing on regional stereotypes and the range of their associated stereotypes. Their study demonstrates that generated pictures of historically marginalized groups exhibit a stereotypical appearance, even when explicitly prompted otherwise. They observed that pictures representing identity groups from countries in South America, Southeast Asia and Africa were the most offensive. Bianci *et. al* (2023) got similar results with Stable Diffusion, noting that harmful stereotypes persist despite mitigation strategies and in some cases, the model amplified stereotypes. They noted that even prompts that don't reference race, gender or ethnicity lead the model to reproduce harmful stereotypes. The generated pictures exoticize people with darker skin tones, tie emotionality to stereotypically white feminine features and generate darker-skinned faces with the prompts *'illegal person'* and *'terrorist'*.

Fraser, Kiritchenko and Nejadgholi (2023) compared Midjourney, Stable Diffusion and DALL·E 2, revealing that all the models showed age bias by generating more images of young people and they all had a noticeable lack of representation of darker-skinned women. Like the study by Bianchi *et. al* (2023), the models generated stereotypical content even when the prompts didn't reference race, gender, or ethnicity. However, the trends were inconsistent between the different models, and the results must be considered preliminary due to the small sample size used in the study.

2.3 Gender gap in the AI industry

Women remain underrepresented in the science, technology, engineering, and mathematics (STEM) workforce. Even among women who obtain STEM degrees, there is a lower likelihood of them pursuing careers in STEM compared to their male counterparts. The computer science and math workforce is dominated by men, as women comprise only 27% of the workforce. Women are especially under-represented in engineering, as only about one out of seven engineers are women. (Beede *et. al*. 2011.)

The gender gap in STEM fields affects the AI industry. The gender gap in AI development and research is particularly pronounced. Element AI and WIRED focused on the 4000 researchers that were published at the leading machine learning conferences, NIPS, ICML and ICLR and concluded that only 12 percent of major machine learning researchers are women (Mantha and Hudson 2018). WIRED reviewed Google's AI research papers, where only 10 percent of people listed as working on "machine intelligence" were women. They also reviewed Facebook's AI research group, which listed 115 people, of whom 15 percent were women (Simonite 2018). Other studies have gotten similar results, finding that only 14% of AI research authors are women (Stathoulopoulos, Mateos-Garcia and Owen 2019).

Despite a more generous estimate by the World Economic Forum in 2023, suggesting that approximately 30% of people working in AI are women, the industry remains overwhelmingly male-dominated.

In a study conducted by Schulenberg *et al.* (2023), eight women with experience in different areas of AI design were interviewed to understand their perceptions of gender and its impact on their approach to AI design and development. Some of the participants revealed they feel accustomed to being the only woman in the room and some felt that men had expectations for them to be more nurturing and caring than their male counterparts. The gender imbalance and divide led to participants feeling imposter syndrome and participants recognized that they had to be more assertive to ensure that their voices were heard in the workplace. Many of the women brought up the existence of 'bro code', where men's voices are prioritized by other men and women's contributions are undermined. Moreover, the participants discussed approaches on how to reduce gender biases within the industry and advocated for a more inclusive approach to AI design. The majority of the participants agreed that AI should have a gender-neutral design. They also believed that the current state of AI design is far too gendered, possibly not appealing to the non-male users of the tools. They supported designing AI tools with use-case specificity and intersectional practises in mind. (Schulenberg *et al.* 2023)

The under-representation of women in the AI workforce creates notable challenges to the field. The aforementioned study by Schulenberg *et al.* (2023) highlights how women find themselves to be in the minority, face stereotypical expectations, and feel that they must modify their behaviour to combat marginalization. According to the female AI design professionals in the study by Schulenberg *et al.* (2023), AI design should focus on an intersectional and user-specific design process to reduce gender bias in the development of new AI technologies. To make applications and research responsive to a broader and diverse user base, a study by Schiebinger and Klinge (2013) advocates for gender analysis to enhance the lives of humankind. Their study concludes that integrating gender analysis into research promotes creativity by providing new perspectives, raising new questions, and opening new research areas. They advocate for using methods of gender analysis to help in rethinking stereotypes and opening design to more dynamic representations of gender. An effort towards systematic change within the industry is needed to address these issues to foster an environment in the AI workforce which promotes diversity and inclusivity.

3 Research process

The following chapter explains how the occupations were selected for the study. Subsequently, DALL·E 3 and Ideogram were prompted with these occupations. When applicable, occupations were additionally prompted with 'student' and 'professor' (for example, ‘Chemistry student’) for further research to uncover potential biases. The collected image data were analysed using thematic analysis.

3.1 Data collection

The occupations were sourced from Statistics Finland’s *‘Employed persons by occupation group (Classification of Occupations 2010, levels 1 to 5)’* from the year 2020 (Statistics Finland, 2023). The classification system used, known as the National Classification of Occupations 2010 is based on ISCO-08 (International Standard Classification of Occupations), developed by the International Labour Organization (Lehto and Ylönen 2011). Occupations were selected from all ten main groups, and each chosen occupation had at least 300 workers or entrepreneurs per occupation in Finland. Gender distribution was a key component in occupation selection. The occupations were divided into five groups based on the number of women workers: ‘Women’s occupations’ (Over 90% of the workforce are women), ‘Women-dominated occupations’ (60% – 90% of the workforce are women), ‘Gender-balanced occupations’ (40% – 60% of the workforce are women), ‘Man-dominated occupations’ (10 – 40% of the workforce are women) and ‘Men’s occupations’ (<10% workforce consists of women). This study includes 15-20 occupations in each category to represent a range of occupations with varying gender ratios across different fields.

DALL·E 3 and Ideogram were prompted to generate images based on the chosen occupations to investigate what kind of stereotyped images AI models produce within a professional context, and to identify differences between the image generators. For each occupation, both DALL·E 3 and Ideogram generated four images, resulting in a total of 680 images. Images that did not show the workers' faces or were otherwise deemed unusable were discarded. For example, some images with the prompts ‘gardener’ and ‘firefighter’ had a helmet or a hat covering the face. In some images, the worker was depicted from an angle that did not show the face. In rare instances, the images portrayed something different compared to what the prompt was (for example, a fantasy character or a robot) and were therefore discarded. Out of the 680 pictures, a total of 602 were used for thematic analysis, from which 291 were generated by Ideogram and 311 by DALL·E 3. Additionally, certain occupations were further analysed by prompting with 'student' and 'professor' to explore potential biases. This approach generated an additional 176 images across eleven occupations. Out of the 176 images, a total of 167 were used, from which 81 were generated by Ideogram and 86 were generated by DALL·E 3. A closer look at the image collection process is provided in Figure 1. Images generated by Ideogram were generated in March of 2024 using Ideogram 1.0 and images generated by DALL·E 3 were generated in March and April of 2024 using ChatGPT-4.

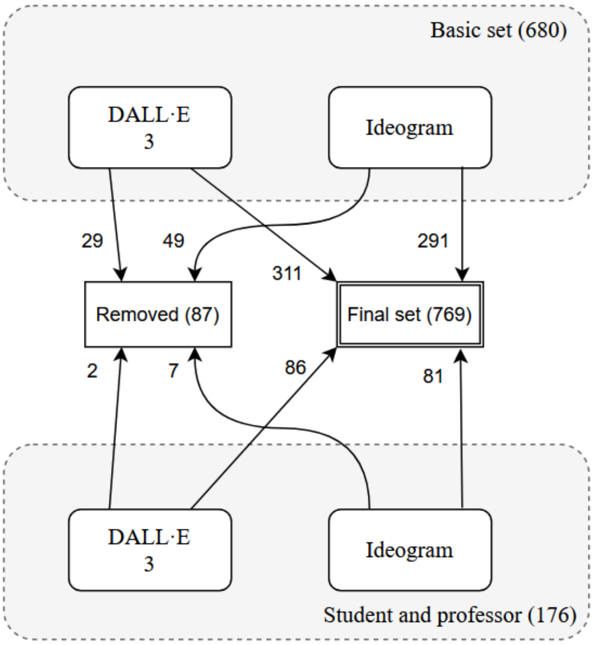

Figure 1: Image collection process

3.2 Data analysis

The images generated by DALL·E 3 and Ideogram were analysed using thematic analysis as per Braun and Clarke (2006). This analysis method involves identifying, analysing, and reporting themes (patterns) in data, organizing and describing the dataset in rich detail, often with further interpretations of various aspects of the research topic. (Boyatzis 1998, cited in Braun and Clarke 2006). Thematic analysis was applied as it is a notably flexible analysis method. Thematic analysis is suited for a wide range of data types; it can be used to analyse textual data from qualitative surveys, story-based methods, online discussion forums, diaries, interviews, and other media sources (Willig and Rogers 2017). Thematic analysis has been efficiently used to analyse images across multiple studies (for example in Del Gesso 2021; Shanahan, Brennan and House 2019).

The six phases of thematic analysis as outlined by Braun and Clarke (2006) were conducted. The phases can be seen in Figure 2. Initially, every data item (generated image) was studied to collect initial ideas and identify possible patterns in the dataset. The notes from this phase were used in phase two to generate initial codes to organize the data. The initial codes fell into four categories: gender, expression, focus of gaze, signs of ageing and clothing and accessories. These initial codes serve as a basis for comparing DALL·E 3 and Ideogram in the following chapter.

Each image received full and equal attention, often receiving multiple codes. In the third phase, the codes were organized into potential themes. An early thematic map was created during this state, which produced a collection of possible themes and sub-themes. In phase four, certain initial themes were either discarded or merged with other themes. For example, themes regarding the focus of gaze were removed, as there was no evident data to support that there were differences across genders regarding gaze. Similar adjustments were made regarding certain clothing elements and other initial codes. After refining the thematic map, the full data set was re-examined to ensure that the map reflects the data set as a whole and to ensure that possible themes were not missed in earlier stages. The themes and sub-themes were defined and named, and a detailed analysis of the themes is presented in the next chapter.

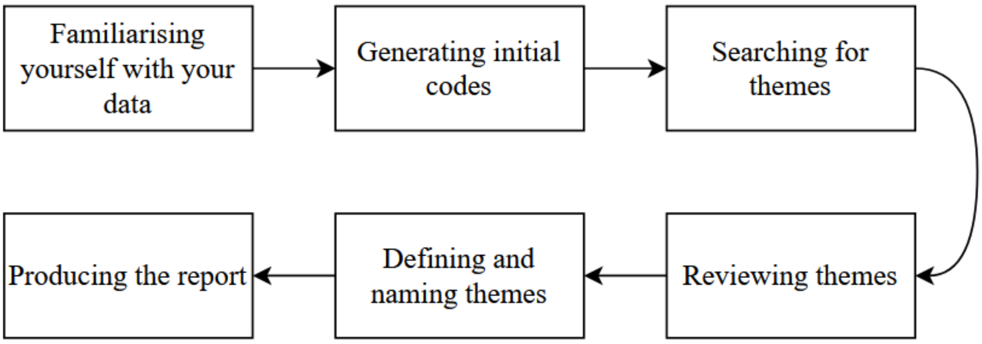

Figure 2: Six phases of thematic analysis as per Braun and Clarke (2006)

4 Results

The results from the thematic analysis are presented in this chapter. The themes identified in the generated images are categorized by gender: ’Warm female persona’ and ’The analytical male persona’. The final thematic map can be seen in Figure 3. Additionally, the results of the student and professor analysis are shown. Finally, a comparison of the gender gap between the images generated by DALL·E 3 and Ideogram is presented.

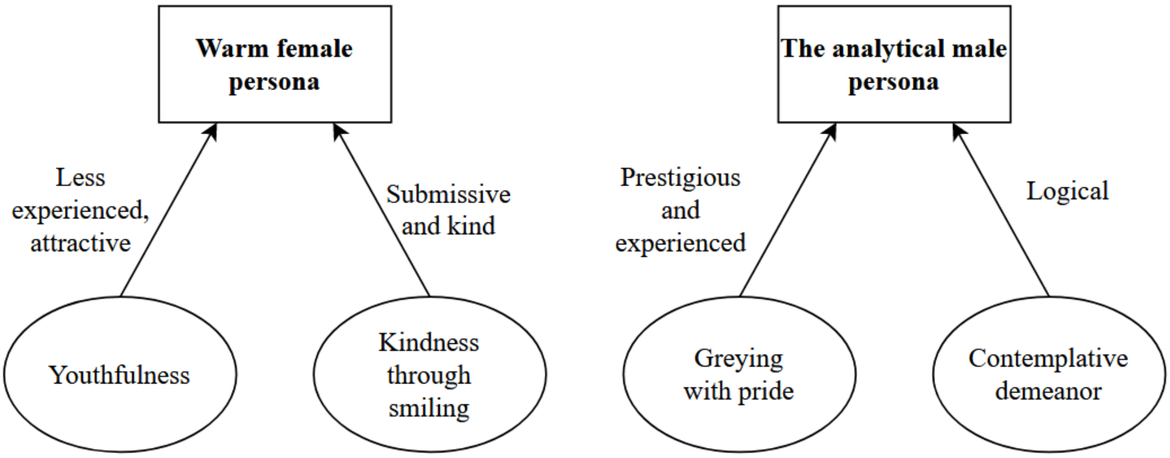

Figure 3: Final thematic map

4.1 Warm female persona

The women in AI-generated images tend to highlight traditionally stereotypical female characteristics, that make the women seem more friendly, co-operative, and submissive in a work setting. This depiction was found in images generated by both models. The theme is divided into two sub-themes that manifest the finding: (a) youthfulness and (b) kindness through smiling.

4.1.1 Kindness through smiling

In the pictures generated by both DALL·E 3 and Ideogram, women were significantly more likely to display positive emotions in an occupational setting. An image of a smiling girl or a woman can be considered as a stereotype for women to act more "lady-like" or well-behaved compared to men (Rodgers, Kenix and Thorson 2007). In a study by Rodgers, Kenix and Thorson (2007), news photographs were examined to find out how individuals of different ages, gender and ethnic subgroups were portrayed. Their study found that on average, a picture of a woman in a newspaper was smiling, had a calm demeanour and a submissive stance. The same phenomenon can be found in AI-generated images generated by DALL·E 3 and Ideogram. Figure 4 shows dental therapists generated by DALL·E 3. DALL·E 3 generated one woman and three men. The woman has a beaming smile and is looking towards the camera, with her hands clasped in front of her. The men are depicted with neutral expressions. One of them is operating on a patient, and the others have their arms crossed. The woman's body language conveys an inviting and kind presence, while the men appear more serious and professional.

Figure 4: Dental therapists generated by DALL·E 3

The same occurrence can be observed in images generated by Ideogram. Figure 5 shows Ideograms depiction of specialist medical practitioners. Ideogram generated two women and two men. The women smile widely, making eye contact with the viewer. The men are focused on their job, working with different technologies.

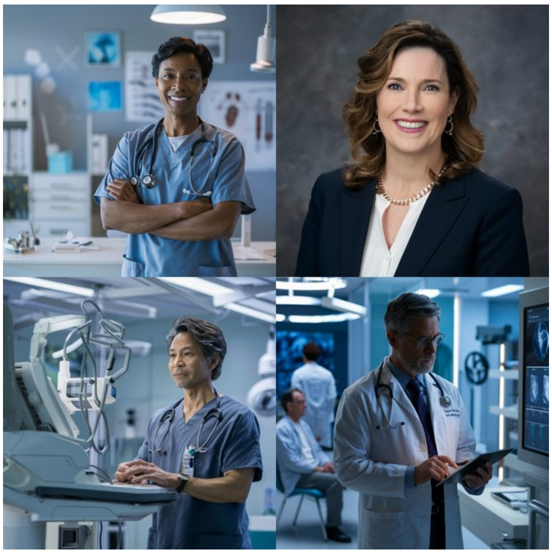

Figure 5: Specialist medical practitioners generated by Ideogram

Table 1 shows the difference in the expressions between genders in pictures generated by DALL·E 3 and Ideogram. In images generated by Ideogram, 74,5% of the depicted women are smiling, compared to 33,3% of smiling men. In comparison, images generated by DALL·E 3 also show a similar trend but with fewer smiling individuals overall: 66,7% of women are smiling, compared to 27,6% of men smiling. When smiling, women were much more likely to smile widely, with their teeth showing in images generated by Ideogram, though both genders had a wide smile more often than a more neutral one. In images generated by DALL·E 3, both genders were more likely to be more neutral with their smile, though women were more likely to smile widely compared to men. In total, the models created 602 pictures, with 192 featuring women, and out of these 72,9% are smiling. The remaining 410 pictures were of men, and out of them, 29,5% were smiling.

Table 1: Expressions of women and men in images generated by DALL·E 3 and Ideogram

Ideogram	**Women**	**Men**	**%/W**	**%/M**	**DALL·E 3**	**Women**	**Men**	**%/W**	**%/M**
Toothy smile	93	31	60,78 %	22,46 %	Toothy smile	12	21	30,77 %	7,72 %
Smile	21	15	13,73 %	10,87 %	Smile	14	54	35,90 %	19,85 %
Neutral	14	34	9,2 %	24,6 %	Neutral	4	61	10,3 %	22,4 %
Focused	24	54	15,7 %	39,1 %	Focused	9	134	23,1 %	49,3 %
Serious	1	4	0,7 %	2,9 %	Serious	0	2	0,0 %	0,7 %
Total	153	138	100 %	100 %	Total	39	272	100 %	100 %

4.1.2 Youthfulness

The portrayal of women as young and attractive in media is a well-studied phenomenon. Women’s magazines are one of the more obvious examples of this, where women's beauty is characterized by (mostly) white skin, thinness, and youth (Byerly and Ross 2008). Youthfulness in an occupational setting is often associated with less experience and authority, while age is respected in many cultures on its own. A long career is respected for the several benefits that come with experience. A worker with a long work history has a track record to show their commitment and reliability in the workplace. Additionally, employees with a long career in a field have been able to acquire a lot of institutional knowledge and high levels of skill in their respective fields.

However, the images generated by DALL·E 3 and Ideogram continue the trend of portraying women as younger than their male counterparts in a professional setting. For instance, when prompting Ideogram with ’physicist’, it generated one picture of a woman and three men as seen in Figure 6. The men depicted show clear signs of ageing: Greying hair and beard, with noticeable wrinkles. In contrast, the female physicist is portrayed with clear and smooth skin, luscious hair, big earrings, and people looking over her shoulder in the background. While it is positive to depict a younger woman in a STEM field, portraying the woman as young and beautiful, compared to the men being old reinforces traditional stereotypes on how women and their worth is tied to their looks and youth.

Figure 6: 'Physicist' generated by Ideogram

When portraying older women, signs of greying were almost unheard of. Out of the 192 images of women, only four had greying hair. When depicting older women, there was high effort put into looks. Older women generated by Ideogram were more likely to have accessories, like earrings and necklaces, visible makeup and curled hair compared to younger women to achieve a well-put-together look. Many of the images of older women overlapped with the subtheme 'Kindness through smiling'. This is demonstrated in Figure 7, which shows women generated by Ideogram with prompts *'Dental therapist'*, *'Nutritionist'* and *'Child care services manager'*.

Figure 7: Older women generated by Ideogram.

In the images generated by Ideogram, older professionals were represented though older women were distinctly underrepresented compared to men. To analyse ageing, signs of ageing were divided into two groups: Wrinkles and greying (hair, beard, or both). In the images generated by Ideogram, over 44% of men showed some signs of ageing, while only 19% of women showed signs of ageing, showing a clear age bias between women and men.

Images generated by DALL·E 3 showed fewer signs of ageing overall. Among these, only one woman depicted by DALL·E 3 showed any signs of ageing, making older women nearly invisible in an occupational context. Wrinkles were an especially rare sight in the pictures generated by DALL·E 3, and most people regardless of gender were portrayed with smooth, unblemished skin. Approximately one in five men had signs of ageing in these images. In total, 15,6% of images depicting women generated by both AI models showed signs of ageing compared to 29,8% of men, showing a significant difference between genders. Table 2 gives a more detailed breakdown of signs of ageing in the images generated by Ideogram and DALL·E 3.

Table 2: Signs of ageing in images generated by Ideogram and DALL·E 3.

Ideogram	Women	Men	%/W	%/M	DALL·E 3	Women	Men	%/W	%/M
Wrinkles	28	58	18,3 %	42,0 %	Wrinkles	1	24	2,6 %	8,8 %
Greying	4	40	2,6 %	29,0 %	Greying	0	48	0,0 %	17,6 %
Wrinkles or greying	29	62	19,0 %	44,9 %	Wrinkles or greying	1	60	2,6 %	22,1 %

4.2 Analytical male persona

The men generated by AI models are often depicted as older and expressing emotions in a more neutral manner. Portraying men in this way makes men seem more experienced and contemplative, making them appear more analytical and experienced compared to women. The theme is divided into two subthemes: (a) greying with pride, and (b) contemplative demeanour.

4.2.1 Greying with pride

Contrary to images of women where signs of greying were almost non-existent, the models effortlessly generated images of older men. In images generated by Ideogram, 29% of men have greying hair or beards, compared to 17.6% of men in images generated by DALL-E 3 (Table 2). Older people, especially women, are underrepresented in media. Research such as Edström (2018) shows that age and gender intersect with media coverage. The study by Edström found that in Swedish media, women are more likely to be visible the younger they are, compared to men whose prime time in media is in their thirties and early forties. Additionally, women who are over the age of 45 were less showcased in the media compared to men. The same phenomena can be found in AI-generated images, where the representational bias reflects the biases in the broader media landscape. There is evidence that age is associated differently between genders in an occupational setting, favouring older men over women. For instance, a study by García-Mainar, Montuenga and García-Martín (2018) found that men receive more prestige as they age, while women receive the same amount or less. Consequently, both models portray men as older, which could continue the trend of representing men as more experienced and prestigious in professional settings.

Figure 8 shows Ideogram's and DALL·E 3's depictions of older men. The top row (geologist, specialist medical practitioner, economist) is generated by DALL·E 3, and the bottom row (Building caretaker, locomotive engineer, construction manager) is generated by Ideogram. The models differ significantly in their representation of age. DALL·E 3 tends to show age by greying hair, omitting wrinkles, making the men seem younger compared to the men depicted by Ideogram. Men generated by DALL·E 3 often conform to conventional media beauty standards, by having a chiselled jawline, thick eyebrows, and a more youthful appearance. Thick eyebrows and a well-defined jawline tend to be perceived as masculine and attractive in male faces (Mogilski and Welling, 2018). In contrast, the men depicted by Ideogram are more varied and realistic by showing clear signs of ageing. While the models have a different way of showing ageing, both feature men with greying hair substantially more than women.

Figure 8: Older men by DALL·E 3 (top row) and Ideogram (bottom row)

4.2.2 Contemplative demeanour

While over half of the women generated by the models are smiling, over half of the men are depicted with a neutral or focused expression. The emotions that women are stereotypically expected to show in the workplace are associated with powerlessness, such as compassion, which contradicts the emotions expected to be displayed by leaders, such as pride (Shields 2010, cited in Taylor *et al.* 2022). One explanation for the difference between genders is that women are more likely to work in an occupation that is higher in emotional labour demands, like healthcare, which emphasizes nurturing skills (Taylor *et al.* 2022). To take the occupational differences in emotional labour better into account, the number of smiling men and women in occupations that have a higher number of women are compared in Tables 3 and 4. Smiling is indeed more apparent in these occupations in images generated by both models, though women are still more likely to smile compared to men. As shown in Table 3, in images generated by Ideogram, the number of women smiling in occupations where the majority are women is ~82%, compared to ~74% in all occupations, and men jump from ~33% to 50%.

Table 3: Quantity of smiling workers by gender, Ideogram

	Ideogram	**Women**	**Men**	**%/W**	**%/M**
Women-dominated occupations	Smile	6	3	13,04 %	12 %
	Toothy smile	32	9	69,57 %	36 %
Women's occupations	Smile	7	1	12,50 %	100 %
	Toothy smile	39	0	69,64 %	0 %
	Total	84	13	82,35 %	50,00 %

In images generated by DALL·E 3, the number of women smiling goes from ~67% to ~72%, and men from ~28% to ~40% as shown in Table 4.

Table 4: Quantity of smiling workers by gender, DALL·E 3

	DALL·E 3	**Women**	**Men**	**%/W**	**%/M**
Women-dominated occupations	Smile	0	18	0,00 %	25,00 %
	Toothy smile	2	9	28,57 %	12,50 %
Women's occupations	Smile	14	9	48,28 %	27 %
	Toothy smile	10	6	34,48 %	0 %
	Total	26	42	72,22 %	40,00 %

The stereotype of women being more emotional, as well as emotionally expressive (e.g., smiling) is well-documented (Lewis, Haviland-Jones and Barrett 2008). This stereotype may influence men to not show their emotions, especially in a professional setting where violating these gender stereotypes may lead to rejection and discrimination. In some studies, men report expressing pride and confidence more frequently compared to women (Lewis, Haviland-Jones and Barrett 2008). These qualities are expected from leaders and as such are often viewed positively in a professional context, reinforcing gender roles in workplaces. Consequently, women showing these traits, that are stereotypically more masculine may face negative social consequences. A thoughtful expression portrays the men as logical and analytical in the workplace, while the women are shown as more submissive and kind.

4.3 Student and professor analysis

The student and professor analysis was done on eleven occupations. The differences between Ideogram and DALL·E 3 can be seen in Table 5. Both models generated professors as older, and students as younger. The biggest difference between the models is the gender gap. DALL·E 3 generated considerably fewer women professors and students compared to Ideogram (Table 5). Ideogram generated a more balanced gender ratio, though most professors are portrayed as men while most students are portrayed as women. In the images, men have a higher academic status, and women are shown as less experienced learners. On the contrary, portraying more students as women may be because nowadays women are increasingly enrolling into higher education.

Table 5: Gender ratio and signs of ageing in ‘professor’ and ‘student’ images.

Ideogram	Women	Men	Wrinkles	Greying	DALL·E 3	Women	Men	Wrinkles	Greying
Professor	13 (30%)	30 (70%)	38 (88%)	34 (79%)	Professor	1 (2%)	42 (98%)	26 (60%)	33 (77%)
Student	21 (55%)	17 (45%)	0	0	Student	1 (2%)	42 (98%)	0	2 (5%)
Total	34 (42%)	47 (58%)	38 (47%)	34 (42%)	Total	2 (2%)	84 (98%)	26 (30%)	35 (41%)

4.4 Gender representation between the models

The occupations were divided into five groups based on the share of women as described in Chapter 3.1. to compare the gender representation between the models. The occupations are divided by the occupational statistics from Finland, and some professions may fall into a different category depending on the country. The ‘student’ and ‘professor’ images are left out of this comparison, as they were analysed in the previous chapter.

Ideogram achieved a balanced representation of women and men in an occupational setting. Four out of five categories fell into the real-life distributions, and women were slightly overrepresented in ‘Man-dominated’ occupations compared to Finnish statistics (34,78% instead of 10% – 30%). Overall, 153 images out of 291 images (52,58%) portrayed women, representing both genders fairly in an occupational context. Figure 9 shows the gender representation across occupational groups in images generated by Ideogram. The number of women in each category is labelled on top of the columns.

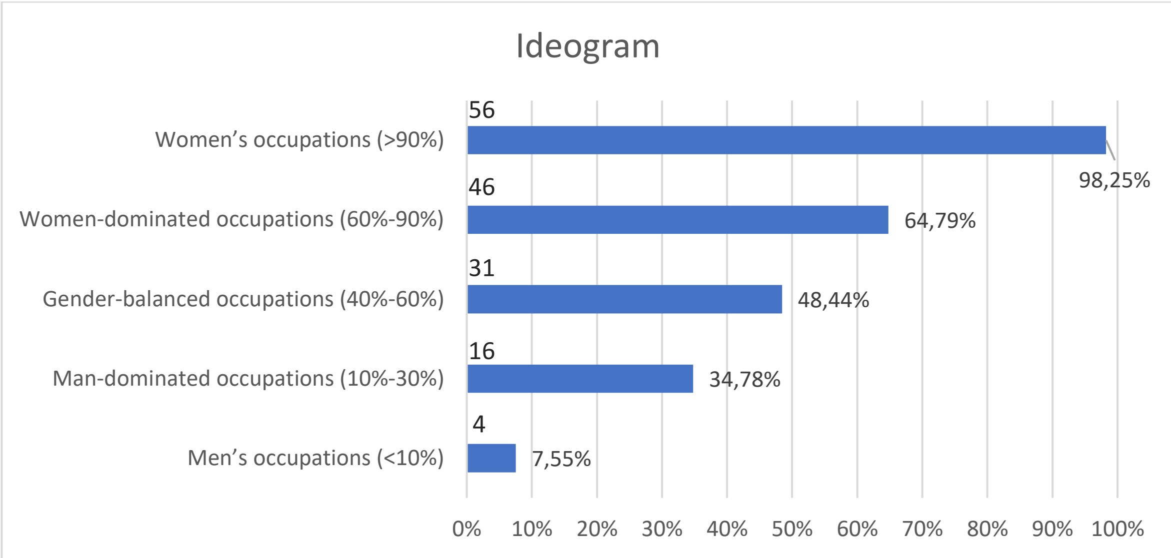

Figure 9: Gender ratio in images generated by Ideogram across occupational groups.

DALL·E 3 exceedingly overrepresented men in an occupational context across all five groups. Women were almost invisible in occupations dominated by men and in occupations that are gender balanced. Even when generating pictures of occupations where over 90% of the workforce are women, less than half of the images generated by DALL·E 3 portrayed women. A closer look at the gender ratio is provided in Figure 10.

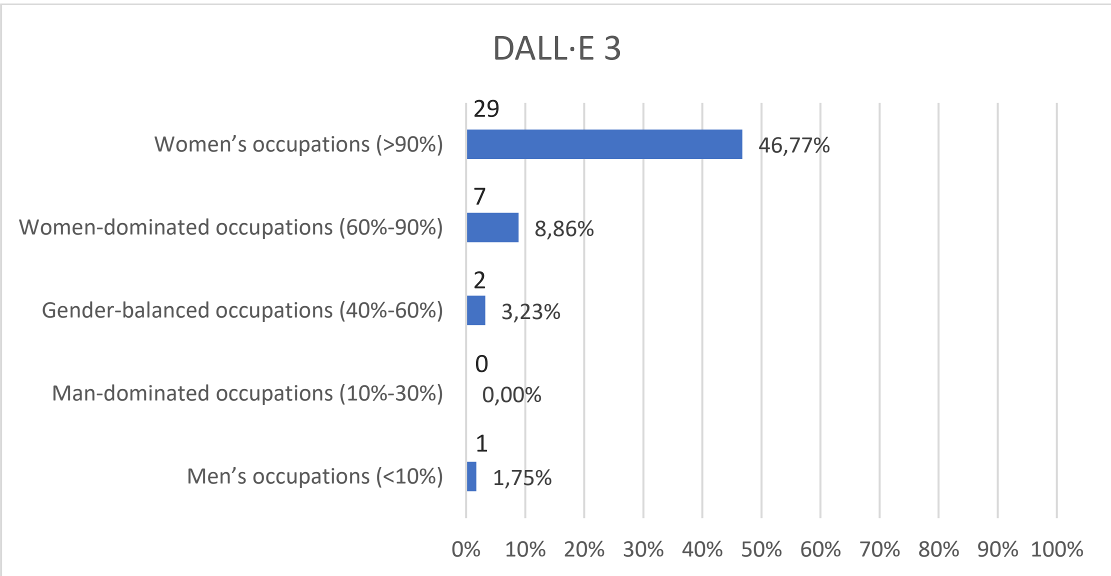

Figure 10: Gender ratio in images generated by DALL·E 3 across occupational groups.

5 Discussion

In this bachelor's thesis, DALL·E 3 and Ideogram were used to explore the presence of representational gender biases in AI-generated pictures in an occupational setting. The results of the analysis answered the research questions: Q1: What kind of stereotyped images do AI models produce within a professional context? and Q2: What noticeable differences in gender bias there are between images generated by DALL·E 3 and Ideogram?

The models differ notably in the number of women depicted in AI-generated pictures. A possible explanation for this difference could be that OpenAI has significantly reduced the number of images DALL·E 3 generates from a single prompt. Ideogram generated four images per prompt with more diversity. In comparison, DALL·E 3 had to be prompted multiple times to generate four pictures. Related studies have shown that DALL·E 2 underrepresents women in an occupational setting and with prompts that are visually under-specified (Sun *et al*. 2023; Fraser, Kiritchenko and Nejadgholi 2023). This study shows an even greater underrepresentation of women in an occupational setting compared to the findings of Sun *et al.* (2023). Men represented 68% of the images in this study, aligning closely with the study of Gorska and Jemielniak (2023), where men represented 76% of professional images across nine image tools. Ideogram had a balanced gender ratio that reflected the real-life gender distribution in occupations, suggesting that it is possible to generate realistic and representative gender ratios.

The results from the thematic analysis align with related studies. DALL·E 3 and Ideogram portray women more than men with smiling faces in an occupational setting, similar to the study by Sun *et al.* (2023) with DALL·E 2. Thematic analysis revealed that AI-generated images depict women stereotypically. Women were younger and smiled more compared to men who were typically older and more neutral with their demeanour in a work setting. DALL·E 3 showed fewer signs of ageing in images of both genders compared to Ideogram, though both models depicted men as older than women. The previous version, DALL·E 2 was shown by Fraser, Kiritchenko and Nejadgholi (2023) to prefer younger age and the results indicate that this bias persists in DALL·E 3. This study contributes to existing literature, confirming that stereotypical gender depictions are prevalent in AI-generated images.

The study has some limitations. It does not consider factors like race or body weight and focuses mostly on gender stereotypes. A more intersectional approach could provide a deeper understanding of the underlying stereotypes in AI-generated images. Gender is classified into two categories (men/women) in the study to simplify comparison with occupational statistics. Future studies could incorporate other gender identities into the research. Finally, as the technology around text-to-image models is advancing rapidly, the results of this study only reflect the current state of these models and are likely to change in the future.

6 Conclusions

In the last few decades, technology has advanced significantly. Simultaneously, harmful gender stereotypes have been observed even in the most popular applications. As text-to-image models have seen rapid advancements in recent years, it is increasingly important to assess what kind of stereotyped images these AI models produce.

This study compared DALL·E 3 and Ideogram to uncover gender biases in images generated by AI models. The primary object was to investigate the presence of representational gender biases in AI-generated pictures in an occupational setting and compare these AI models to one another. Over 750 images were generated of occupations with varying gender ratios across different fields. The images were then analysed by using thematic analysis.

The results of the thematic analysis revealed that both models represent women as younger and more likely to smile than men. Over half of the women generated by the models were smiling, while over half of the men were portrayed with a neutral or focused expression. Ideogram showed a higher representation of older professionals, though both models underrepresented older women. When depicting older women, greying hair was rarely seen. Ideogram had a balanced representation of women and men in a professional setting, while DALL·E 3 considerably overrepresented men in all occupational groups.

The results indicate that AI models reinforce gender roles by portraying women as younger and more likely to smile than their male counterparts. Men were more likely to be portrayed as older in an occupational context, possibly depicting them as more experienced and professional. DALL·E 3 and Ideogram underrepresent older women, which is a known phenomenon in media.

The findings show that continuously evaluating new and existing text-to-image generators is important. Future studies could reveal if the stereotypical portrayal of genders continues. Increasingly more images in media are AI-generated, and future research could examine the possible effects of AI-generated imagery on a larger scale.

References

Beede, D.N., Julian, T.A., Langdon, D., McKittrick, G., Khan, B. and Doms, M.E. (2011) 'Women in STEM: A Gender Gap to Innovation', *Economics and Statistics Administration Issue Brief*, 04-11, Rochester, NY. Available at: https://doi.org/10.2139/ssrn.1964782.

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 *2023 ACM Conference on Fairness, Accountability, and Transparency*, pp. 1493–1504. Available at: https://doi.org/10.1145/3593013.3594095.

Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', *Qualitative Research in Psychology*, 3, pp. 77–101. Available at: https://doi.org/10.1191/1478088706qp063oa.

Byerly, C.M. and Ross, K. (2008) *Women and Media: A Critical Introduction*. John Wiley & Sons. p.49. ISBN: 978-1405116077.

Cave, S. and Dihal, K. (2020) 'The Whiteness of AI', *Philosophy & Technology*, 33(4), pp. 685–703. Available at: https://doi.org/10.1007/s13347-020-00415-6.

Chang, Z., Koulieris, G.A. and Shum, H.P.H. (2023) 'On the Design Fundamentals of Diffusion Models: A Survey'. arXiv. Available at: https://doi.org/10.48550/arXiv.2306.04542.

Del Gesso, C. (2021) 'A picture is worth a thousand words: a photo-thematic analysis of city hospitality in municipal popular reporting', *Journal of Hospitality and Tourism Technology*, 13(1), pp. 100–119. Available at: https://doi.org/10.1108/JHTT-08-2020-0206.

Dhariwal, P. and Nichol, A. (2021) 'Diffusion Models Beat GANs on Image Synthesis', in *Advances in Neural Information Processing Systems*. Curran Associates, Inc., pp. 8780–8794. Available at: https://proceedings.neurips.cc/paper_files/paper/2021/hash/49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.

Edström, M. (2018) 'Visibility patterns of gendered ageism in the media buzz: a study of the representation of gender and age over three decades', *Feminist Media Studies*, 18, pp. 77–93. Available at: https://doi.org/10.1080/14680777.2018.1409989.

Ellemers, N. (2018) 'Gender Stereotypes', *Annual Review of Psychology*, 69, pp. 275–298. Available at: https://doi.org/10.1146/annurev-psych-122216-011719.

Fessler, L. (2017) *We tested bots like Siri and Alexa to see who would stand up to sexual harassment*, *Quartz*. Available at: https://qz.com/911681/we-tested-apples-siri-amazonechos-alexa-microsofts-cortana-and-googles-google-home-to-see-which-personal-assistant-bots-stand-up-for-themselves-in-the-face-of-sexual-harassment (Accessed: 14 April 2024).

Fraser, K.C., Kiritchenko, S. and Nejadgholi, I. (2023) ‘A Friendly Face: Do Text-to-Image Systems Rely on Stereotypes when the Input is Under-Specified?’ arXiv. Available at: https://doi.org/10.48550/arXiv.2302.07159.

García-Mainar, I., Montuenga, V. and García-Martín, G. (2018). ‘Occupational Prestige and Gender-Occupational Segregation’, *Work, Employment and Society,* 32(2), pp. 348-367. Available at: https://doi.org/10.1177/0950017017730528.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. (2020) ‘Generative adversarial networks’, *Communications of the ACM*, 63(11), pp. 139–144. Available at: https://doi.org/10.1145/3422622.

Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y. (2014) ‘Generative Adversarial Networks’. arXiv. Available at: https://doi.org/10.48550/arXiv.1406.2661.

Gorska, A.M. and Jemielniak, D. (2023) ‘The invisible women: uncovering gender bias in AI-generated images of professionals’, *Feminist Media Studies*, 23(8), pp. 4370–4375. Available at: https://doi.org/10.1080/14680777.2023.2263659.

Hu, K. (2023). ChatGPT sets record for fastest-growing user base - analyst note. Available at: https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/(Accessed: 27.01.2024).

Ideogram. (2023). Ideogram v0.1 is open to everyone for free. Available at: https://ideogram.ai/publicly-available (Accessed: 27.01.2024).

Jha, A., Prabhakaran, V., Denton, R., Laszlo, S., Dave, S., Qadri, R., Reddy, C.K. and Dev, S. (2024) ‘ViSAGe: A Global-Scale Analysis of Visual Stereotypes in Text-to-Image Generation’. arXiv. Available at: https://doi.org/10.48550/arXiv.2401.06310.

Kay, M., Matuszek, C. and Munson, S.A. (2015) ‘Unequal Representation and Gender Stereotypes in Image Search Results for Occupations’, in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. New York, NY, USA: Association for Computing Machinery (CHI ’15), pp. 3819–3828. Available at: https://doi.org/10.1145/2702123.2702520.

Kuczmarski, J. (2018). Reducing gender bias in Google Translate. Available at: https://blog.google/products/translate/reducing-gender-bias-google-translate/(Accessed: 03.03.2024).

Lehto, K. and Ylönen, S. (2011). Ammattiluokitus 2010. Available at: https://www.doria.fi/bitstream/handle/10024/103626/yksk14_201000_2011_net.pdf?seque Accessed: 23.04.2024).

Lewis, M., Haviland-Jones, J.M. and Barrett, L.F. (2008) Handbook of Emotions, Third Edition. Guilford Press, pp.396-398. ISBN: 978-1609180447.

Mantha, Y., Hudson, S. (2019). Estimating the gender ratio of AI researchers around the world. Available at: https://medium.com/element-ai-research-lab/estimating-the-gender-ratio-of-ai-researchers-around-the-world-81d2b8dbe9c3 (Accessed: 28.02.2024).

Mishkin, P., Ahmad, L., Brundage, M., Krueger, G. and Sastry, G. (2022). DALL·E 2 Preview - Risks and Limitations. Available at: https://github.com/openai/dalle-2-preview/blob/main/system-card.md (Accessed 27.01.2024).

Mogilski, J.K. and Welling, L.L.M. (2018) 'The Relative Contribution of Jawbone and Cheekbone Prominence, Eyebrow Thickness, Eye Size, and Face Length to Evaluations of Facial Masculinity and Attractiveness: A Conjoint Data-Driven Approach', *Frontiers in Psychology*, 9. Available at: https://doi.org/10.3389/fpsyg.2018.02428.

OpenAI. (2022a). Introducing ChatGPT. Available at: https://openai.com/blog/chatgpt (Accessed: 27.01.2024).

OpenAI. (2022b). DALL·E 2. Available at: https://openai.com/dall-e-2 (Accessed: 27.01.2024).

OpenAI. (2023). DALL·E 3. Available at: https://openai.com/dall-e-3 (Accessed: 27.01.2024).

Ramesh, A., Dhariwal, P., Nichol, A., Chu, C. and Chen, M. (2022) 'Hierarchical Text-Conditional Image Generation with CLIP Latents'. arXiv. Available at: https://doi.org/10.48550/arXiv.2204.06125.

Rodgers, S., Kenix, L.J. and Thorson, E. (2007) 'Stereotypical Portrayals of Emotionality in News Photos', *Mass Communication and Society*, 10(1), pp. 119–138. Available at: https://doi.org/10.1080/15205430709337007.

Rombach, R., Blattmann, A., Lorenz, D., Esser, P. and Ommer, B. (2022) 'High-Resolution Image Synthesis with Latent Diffusion Models', in. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE Computer Society, pp. 10674–10685. Available at: https://doi.org/10.1109/CVPR52688.2022.01042.

Samuel, S. (2022) *A new AI draws delightful and not-so-delightful images*, *Vox*. Available at: https://www.vox.com/future-perfect/23023538/ai-dalle-2-openai-bias-gpt-3-incentives (Accessed: 16 May 2024).

Schiebinger, L. and Klinge, I. (2013) 'Gendered Innovations: How Gender Analysis Contributes to Research'. Publications Office of the European Union. Available at: https://doi.org/10.2777/11868.

Schulenberg, K., Hauptman, A.I., Schlesener, E.A., Watkins, H. and Freeman, G. (2023) '``I Felt Like I Wasn't Really Meant to be There'': Understanding Women's Perceptions of

Gender in Approaching AI Design & Development', in. *Hawaii International Conference on System Sciences*. Available at: https://doi.org/10.24251/HICSS.2023.022.

Shanahan, N., Brennan, C. and House, A. (2019) 'Self-harm and social media: thematic analysis of images posted on three social media sites', *BMJ Open*, 9(2). Available at: https://doi.org/10.1136/bmjopen-2018-027006.

Simonite, T. (2018) 'AI Is the Future—But Where Are the Women?', *Wired*. Available at: https://www.wired.com/story/artificial-intelligence-researchers-gender-imbalance/ (Accessed: 28.2.2024).

Stathoulopoulos, K., Mateos-Garcia, J.C. and Owen, H. (2019) 'Gender Diversity in AI Research', *SSRN Electronic Journal* [Preprint]. Available at: https://doi.org/10.2139/ssrn.3428240.

Statistics Finland. (2023). '115q -- Employed persons by occupation group (Classification of Occupations 2010, levels 1 to 5), occupational status, sex and year, 2010-2021'. Available at:https://pxdata.stat.fi/PxWeb/pxweb/en/StatFin/StatFin__tyokay/statfin_tyokay_pxt_115q.px/ (Accessed: 20.04.2024).

Sun, L., Wei, M., Sun, Y., Suh, Y.J., Shen, L. and Yang, S. (2024) 'Smiling women pitching down: auditing representational and presentational gender biases in image-generative AI', *Journal of Computer-Mediated Communication*, 29. Available at: https://doi.org/10.1093/jcmc/zmad045.

Taylor, C.L., Ivcevic, Z., Moeller, J., Menges, J.I., Reiter-Palmon, R. and Brackett, M.A. (2022) 'Gender and Emotions at Work: Organizational Rank Has Greater Emotional Benefits for Men than Women', *Sex Roles*, 86(1), pp. 127–142. Available at: https://doi.org/10.1007/s11199-021-01256-z.

University of Toronto. (2023). Teaching Assistants' Training Program. Available at: https://tatp.utoronto.ca/resources/equity-diversity-and-inclusion/ (Accessed: 27.02.2024).

Wang, J., Liu, X., Di, Z., Liu, Y. and Wang, X. (2023) 'T2IAT: Measuring Valence and Stereotypical Biases in Text-to-Image Generation', *Findings of the Association for Computational Linguistics: ACL 2023*, Toronto, Canada, Association for Computational Linguistics, pp. 2560–2574. Available at: https://doi.org/10.18653/v1/2023.findings-acl.160.

Willig, C. and Rogers, W.S. (2017) *The SAGE Handbook of Qualitative Research in Psychology*. SAGE. ISBN: 978-1473925212.

World Economic Forum. (2023). Global Gender Gap Report 2023. Available at: https://www.weforum.org/publications/global-gender-gap-report-2023/ (Accessed: 28.02.2024).