

Double Disadvantage: An Interview Study with Women in the AI Industry on the Role of Gender in AI Ethics

Maaïke Harbers¹, Kirsten van der Stelt², Sander 't Sas²

¹Rotterdam University of Applied Sciences, The Netherlands

²Independent researcher

Corresponding author: m.harbers@hr.nl

Abstract

With the rapid rise of Artificial Intelligence (AI), increasing attention has been paid to the ethical implications of AI systems. Much of the existing work on AI ethics focuses primarily on technical characteristics of these systems, while paying less attention to the social contexts in which they are developed and deployed. However, prior research suggests that diversity in teams leads to more ethical decision making. Despite the relevance of diversity for AI ethics, the AI workforce is predominantly male, and there is little research, particularly empirical, on the relation between gender and AI ethics. This paper addresses this gap by presenting an interview study with 19 women working in the Dutch AI industry, exploring their perspectives on the role of gender in AI ethics. Using thematic analysis, we defined four key themes: 1) women enrich AI ethics, 2) ethics is undervalued in the AI sector, 3) women are often taken less seriously in AI workplaces, and 4) homogeneity is reinforced in large parts of the AI sector. Our findings suggest that diversity and inclusion, and AI ethics are strongly interconnected, highlighting the importance of addressing structural inequalities within the AI sector to support more ethically informed AI development and deployment.

1 Introduction

In the last decade, AI has become deeply integrated into many areas of society. Despite the promise of AI applications, they also bring serious ethical and societal challenges. These challenges span a wide range of issues, including discrimination against certain groups by algorithms, reinforcement of existing stereotypes and inequalities, large-scale dissemination of misinformation or hate speech, and the growing environmental impact of AI (Stahl 2021). The risks associated with AI are widely acknowledged, and along with the rise of AI, increasing attention has been paid to the ethical aspects of AI (Jobin et al. 2019; Trotta et al. 2023). However, although AI ethics is a widely discussed and extensively researched topic, not all ethical issues receive the same level of attention. Several studies show that most of

the work on AI ethics focuses on issues like transparency, privacy, accountability and fairness (Hagendorff 2020; Khan et al. 2022; Portegies et al. 2025; Prem, 2023; Shams et al. 2025). Hagendorff (2020) and Portegies et al. (2025) argue that these are often issues that can (partly) be addressed by focusing on technical characteristics of AI systems. In contrast, ethical issues that require a focus on the social context in which AI systems are developed, like inclusion, diversity, and workers' rights, receive much less attention (Hagendorff 2020; Khan et al. 2022; Portegies et al. 2025; Prem 2023; Shams et al. 2025).

Though diversity and inclusion considerations are understudied in AI development and deployment, they are important for the ethical impact of AI applications (Shams et al. 2025; Zowghi and da Rimini, 2023). There are numerous examples of AI systems discriminating against groups of people due to biased AI systems (Ferrer et al. 2021; Noble 2018). Research suggests that diverse teams increase the likelihood of addressing biases in AI systems, as people who are potentially affected by AI bias are more likely to recognize, understand, and try to address it (West et al. 2019; Zowghi and da Rimini 2023). In addition to that, diverse teams bring different perspectives to what concepts like 'fair AI' and 'trustworthy AI' involve, and by accounting for these multiple perspectives, they can reduce the risk of unintended consequences of AI systems (Shams et al. 2025). Research also shows that diversity in teams in general leads to more ethical decision making (DeGrassi et al. 2012). Thus, given the frequent occurrence of bias in AI systems and discrimination due to that, it could be argued that diversity in development teams should be prioritized.

An important group that is currently underrepresented in the AI workforce are women (Roopaei et al. 2021). Though data on the percentage of women in AI professions is hard to collect, as the field is hard to define, estimates range from 11.3% globally (Young et al. 2023) to 35% in OECD coun-

tries (Green and Lamby 2023), making women a clear minority. The absence of female perspectives in technology design often results in technology that is less safe or user-friendly for women (Criado Perez 2019). This issue extends to AI technology; there are numerous studies describing gender bias in AI systems (e.g., Nadeem et al. 2020; Nadeem et al. 2022; O'Connor and Liu 2024; West et al. 2019). However, though several researchers have stressed the importance of relating bias in AI systems to the people who develop these systems (West et al. 2019; Zowghi and da Rimini 2023), there is little research, particularly empirical, on the role of gender in relation to AI ethics and the development of responsible AI (Kronqvist and Rousi 2022).

To address this gap, this paper presents a study on the role of gender in AI ethics in practice. The research question of the study is: *Which role do women play in ethical decision making in the AI industry and how is that affected by being female?* We answer this question through an interview study with 19 women who work in an AI profession in the Dutch private sector. We have not been able to find figures on the percentage of women in AI roles in the Netherlands specifically, but according to a national platform monitoring the labor market in the Dutch tech industry, in 2024, only 11% of ICT sector employees in the Netherlands were female (Platform Talent voor Technologie), showing that women are clearly a minority in the sector. We analyzed the results of the interviews through thematic analysis, resulting in four key themes. The insights gained from this study can help shape AI ethics practices in organizations, as well as diversity and inclusion policies, thereby contributing to the mitigation of harmful effects of AI technology.

The remainder of this paper starts by providing a background on related literature in Section 2. We will then describe the method that we followed in this study in Section 3. Subsequently, we will present the main results of the study in Section 4. Finally, we provide a discussion and a conclusion in Section 5 and 6, respectively.

2 Background

An increasing body of research shows that AI technologies can strengthen human biases and stereotypes, creating disadvantages and discriminatory effects for individuals (Benjamin 2019; Ferrer et al. 2021; Noble 2018). Examples in which women are disadvantaged by AI applications include facial recognition software that performs worse for women, especially Black women (Buolamwini and Gebru 2018), recruitment algorithms that match men more often than women with high-paying jobs (Hunkenschroer and Luetge 2022), large language models like ChatGPT that reinforce female stereotypes (Ray 2023), nudification apps that are mostly targeted and used to virtually undress women (Gibson et al. 2025), and virtual assistants like Alexa and Siri

with female appearances and voices, reinforcing the stereotype of women as subservient and compliant (Strengers and Kennedy 2021). The effects of these AI systems affect different groups of people in different ways. Gender bias in AI often intersects with other categories of bias, such as ethnicity, class, sexual orientation, religion, ability, or age (Ulnicane 2024). Often, the more someone deviates from the dominant norm, the greater the likelihood and severity of negative consequences from AI bias (Ulnicane 2024).

Harmful effects of AI technologies have led to an increased attention to AI ethics (Jobin et al. 2019; Trotta et al. 2023). As mentioned in the introduction, multiple researchers argue that most attention in AI ethics goes to ethical issues for which there are technical solutions, rather than issues that require mostly social or socio-technical changes (Hagendorff 2020; Portegies et al. 2025). In addition to that, much of the research on AI ethics focuses on identifying and studying ethical risks posed by AI systems, rather than studying how to mitigate these risks in real-world settings (Morley et al. 2021; Zhou and Chen 2023). Yet, empirical studies on AI ethics show that social, cultural and organizational factors play an important role in establishing a practice of ethical decision making in companies working on AI. For example, an interview study demonstrated that AI designers and developers regularly encountered value tensions and trade-offs between ethical principles in the development of AI systems, while they had little tools or guidance on how to deal with them (Sanderson et al. 2021). Another empirical study discussing different use cases of responsible AI in practice, suggested that in addition to the availability of ethical tools, ethical deployment of AI applications also requires organizational awareness of AI ethics, knowledge of mitigation strategies, and an organizational governance that supports the ethical development and deployment of AI (Harbers et al. 2025). Thus, though there is a lot of research on AI ethics in general, there is much less work on how ethical decision-making regarding AI technologies works in practice, including the impact of aspects like diversity and inclusion in the workforce and work culture.

In this paper, we focus on the experiences of female professionals in the AI workforce. As argued, since women are more often disadvantaged by AI-systems than men, including female perspectives is particularly important for the ethical development of AI technologies (Roopaei et al. 2021). There is ample research on women in the fields of IT and computer science, fields that AI is usually considered part of. This research shows that, although the earliest programmers were women, computer science has a long history of failing to recognize and acknowledge women's contributions (Abbate 2012; Hicks 2017; Margolis and Fisher 2002). In a recent empirical study among professionals in IT companies, women perceived gender as a factor that negatively influenced their careers, e.g. regarding salary and promotion opportunities, whereas men were less aware of the impact of

their gender (Kronqvist and Rousi 2025). Women in the study were also significantly more likely than men to report experiencing sexual harassment, and they reported higher levels of biased or dismissive behavior (Kronqvist and Rousi 2025). In conclusion, research shows that experiences of women in the IT industry significantly differ from those of men. This suggests that it is worthwhile to examine their perspectives and experiences in relation to AI ethics.

To summarize, there is a lot of work on AI ethics, as well as on the experiences of women working in IT. However, there is little research combining these topics. Kronqvist and Rousi (2022) did a literature review on the terms ‘ethics’, ‘design thinking’ ‘gender’ and ‘AI’ and found that the examination of gender and its role in developing ethical AI applications appears to receive limited attention in the literature. To the best of our knowledge, no major empirical studies have been published on this specific topic since.

3 Methods

The study presented in this paper consists of a thematic analysis (Braune and Clarke 2006) to analyze the results of 19 semi-structured interviews with women who worked with or on AI in the Dutch AI industry. This section describes the recruitment process, the participants, and the interview and data analysis processes. In all steps of this study, we followed the ethical and data management procedures of the first author’s institution.

Recruitment. We recruited participants with a call describing the desired characteristics of the participants. The call targeted (self-identified) women who: work or worked with or on AI applications, such as recommenders, image recognition software, chatbots, expert systems, etc.; work(ed) in various roles, e.g., as programmers, UX designers, auditors of AI applications, or as team leaders of an AI team; are/were employed in the private sector in the Netherlands, this could be self-employed, a small or medium-sized enterprise (SME), or a large company; and who speak Dutch or English. It should be noted that at the time of recruitment (late 2023) the integration of large language models (LLMs) in work processes was less widespread, possibly affecting the understanding of the term ‘AI applications’ in the call.

We used the call to recruit participants through various ways, including direct outreach to professionals in the authors’ own networks, a LinkedIn post by one of the authors, and snowball sampling. The LinkedIn post was viewed over 7,000 times and shared more than 20 times, reaching many professionals beyond the direct networks of the authors. Some of the people who expressed interest in participating in the study were working at a research institute or university. These individuals were excluded from our study, as we were targeting people working in the private sector. All other people who showed interest in participation met the

conditions in the recruitment call and were invited to participate. We stopped recruiting after the 20th interview was scheduled. One participant withdrew from the study after the interviews were conducted. Her data were deleted, and the study thus includes 19 participants.

Participants. To protect the anonymity of the participants and their employers, we provide a general description of the participants rather than detailed individual profiles. The participants had different forms of employment: four worked as freelancers, four were employed at an SME, and eleven were employed at a large commercial organization. Two of the participants worked at the same large company. All other participants worked at unique organizations, constituting a total of 18 companies in the study. All the freelancers had previously worked at a company in an AI role. The four SMEs had AI or data analysis as their core business. In most of the ten large companies, AI was not the core business. Instead, participants worked at an AI or IT department which supported the companies’ activities in another sector, such as consulting, finance, HR, logistics, or e-commerce.

The participants held various positions, including programmer, data analyst, AI researcher, product owner, ethics officer, consultant, and manager. Seniority levels ranged from junior to managing director of a multinational. Seven of the participants held formal leadership roles, and three others had informal leadership responsibilities.

We did not ask the participants explicitly about their nationality or cultural background. Yet the topic came up in several interviews. Three of the interviews were conducted in English due to limited Dutch proficiency. This makes it highly unlikely that these participants grew up in The Netherlands. All other interviews were conducted in Dutch. Of the 16 interviews in Dutch, one participant mentioned having a non-Dutch background, another mentioned having a bicultural background, and a third participant explicitly mentioned having an Islamic background.

Interviews and analysis. The questionnaire for the semi-structured interviews was developed by the authors in multiple test rounds. Our aim was to minimize the information provided to the participants, but we added short leading introductions to introduce or explain the concepts of ‘AI ethics’, ‘gender differences in dealing with AI ethics’, and ‘gender bias in AI systems’ to improve participants’ understanding of the questions. We created a Dutch and an English version of the questionnaire. The interviews were conducted between October 2023 and March 2024 by the authors of this paper. Participants were not interviewed by someone they personally knew. All interviews were conducted online via MS Teams, with audio and video recordings. The interviews lasted about one hour. The recordings were transcribed using Amberscript software and manually reviewed and supplemented. Subsequently, we analyzed the transcripts using Atlas.ti in an iterative process of inductive coding, searching for themes and reviewing themes (Braun

and Clarke 2006; Braun and Clarke 2019; Braun and Clarke 2023). In this process, the authors came together multiple times to define and review themes, until consensus was reached on the final themes. The authors added complementary perspectives to the analysis. The first author has a background in AI and philosophy, and the experience of having worked as a female AI researcher in often male-dominated environments for many years, the second author has a background in science and technology studies (STS) and gender studies, and the third author has a background in cultural studies with a focus on gender. To report the results, we translated the Dutch quotes that we included in this paper to English.

4 Results

In the interviews, we asked participants about the gender balance in their workplaces. This question was hard to answer precisely for many of the participants. Many of them noted that the overall gender ratio in their company differed from that in their specific (technical) department or team, with fewer women in technical teams. Several also mentioned a significant difference between the gender ratio on the work floor and in upper management, with fewer women in upper management. Some of the participants also indicated that they did not belong to a single defined team but collaborated across multiple teams. Despite these considerations, reported gender ratios ranged from 10% women and 90% men to a 50/50 split. As we will get back to later, many of the participants who worked at a company or in a team with a more even gender ratio had put effort into working at such a place. All the participants said that they had experience working in male-majority environments.

Almost all the participants reported encountering ethical issues related to AI in their work. Ethical issues mentioned in the interviews included: bias in data, discrimination by algorithms, algorithmic fairness, transparency and explainability, privacy issues, robustness and safety of AI systems, legislation, and environmental impact. In addition to these concerns related to AI technology, some ethical issues were more social in nature. For example, one participant recalled a situation in which someone who identified themselves as part of the LGBTQ+ community was asked to work on an AI project for a Russian client, despite the known hostility toward LGBTQ+ people in parts of Russian society.

In the remainder of this section we will present the results of our thematic analysis (Braun and Clarke 2006), organized into the following four key themes: women enrich AI ethics, ethics is undervalued in the AI sector, women are often taken less seriously in AI workplaces, and homogeneity is reinforced in large parts of the AI sector.

4.1 Women Enrich AI Ethics

We asked the participants whether they believed their gender or other identity characteristics played a role in how they engaged with AI ethics in their work. While some participants explicitly indicated that they were not sure if their observations reflected a broader pattern or were coincidental, we identified three gender differences in relation to AI ethics that were each mentioned by multiple participants.

First, the results show that most participants in this study believed that **women are more likely to recognize ethical issues**. When asked for it, nearly all participants thought that women were more likely than men to recognize gender bias in AI systems, attributing this to personal experience and social position. Many expanded their answers beyond gender, emphasizing the importance of diverse perspectives in general, as illustrated by these two quotes: *“As a person, you’re always bound by your own frame of reference, your life experiences, and what you’ve been through. That includes your height, skin color, physical ability, and of course your gender. So yes, I’m sure that has an influence.”* (P10) and *“I think it’s because of how you experience the world. You notice different things and have different concerns. That applies not just to women, but to diversity across the board.”* (P13). Generational differences were frequently mentioned as an important characteristic for recognizing ethical issues, next to gender, where younger people were seen as more alert to ethical risks, especially regarding climate impact. This is illustrated by this quote: *“If I look at my younger colleagues in their 20s, they’ll notice everything. If they see bias, they’ll speak up. Among people my age, I think there are still more ingrained clichés, and women are more likely to notice them than men.”* (P1).

Some participants emphasized that awareness, and not just identity, plays an important role: *“I’m very engaged with the topic, so I think about it quickly. But I can imagine I have female colleagues who aren’t thinking about it at all.”* (P4) and *“It largely depends on how aware you are as a person. [...] Once you become aware, you can’t unsee it.”* (P8). This suggests that while women may be more likely to notice ethical issues related to AI, other identity characteristics, awareness and education also play a key role.

Second, many of the participants in this study observed that **women more often work on AI ethics**. Many of the participants mentioned that women were strongly represented in activities related to AI ethics, either as a formal part of their job or informally, e.g., by more frequently raising ethical concerns. This trend was noted within companies, as illustrated by this quote: *“I notice that women are more vocal about it at our company. Quite often, women blog about it on our intranet. [...] It’s not that men aren’t involved, but women seem to bring it up or take the lead more often.”* (P12). Participants also noted it outside of their company, for example, at conferences: *“What always*

strikes me, and I find it quite telling, is that at conferences, most attendees are men. But in the ethics track, women are the ones on stage. And if you look at the audience, the gender ratio is completely different from the rest of the conference.” (P7). One of the participants added that this trend also applied to people of color, i.e., people of color are more often concerned with ethical issues in AI. This gender difference seems to have a relation to the previous one; those who are more likely to recognize ethical issues, may also be more likely to address them.

Third, the majority of participants in this study believed that **women offer additional perspectives to ethical issues**. Most participants felt there were differences in how men and women approached ethical questions. Several of the participants pointed out that women are more likely to question whether a technological (AI) solution is appropriate for the given situation and to consider alternatives. The following participant, for instance, questioned the purchase of new hardware: *“One department at our company focuses on robotics, they wanted to deploy more robots instead of people. My colleague and I asked them not just to look at what’s cool or possible, but whether it was really necessary. Because the materials used in robots and their usage aren’t always ethical. Buying them just for the sake of it didn’t sit right with us. In the end, they bought one instead of ten.”* (P12). Another participant said that she frequently directed others’ attention from the technical aspects of an AI-system to the larger context in which it was applied: *“So I’m mostly on the track of: ‘Yeah, the model is cool, but it’s not the full solution.’ How can you embed it in a broader, more sociotechnical system? That opens up all kinds of possibilities, and it’s a real eye-opener for many people.”* (P13). The same participant also noted that male-dominated teams tended to take more risks: *“In teams with a more masculine culture, things are approached more playfully, like, ‘Let’s just try it and see what happens. We can always pull the plug.’ They’re more willing to experiment. In more female-dominated teams, I notice more risk aversion, people are quicker to say, ‘But what if this or that happens?’”* (P13). These findings suggest that women, more often than men, tend to consider the context in which an AI application is applied and whether it is appropriate in that context, rather than merely thinking about its technical possibilities.

Several participants spoke about how their (different) cultural background gave them additional perspectives on ethical issues. This is illustrated by the following quote: *“Women bring a different perspective in how they see things. Coming from outside the Netherlands, I also bring a more outside-in approach, how things are done in other industries or countries, and that helps.”* (P19). This suggests that not just female perspectives can enrich AI ethics, but that having a perspective that deviates from the norm in general can be valuable for AI ethics.

4.2 Ethics Is Undervalued in the AI Sector

A topic that came up in many of the interviews was the low status of ethics in the field of AI. Many participants said that they spoke up about ethical concerns, but that this often was not met with enthusiasm by their colleagues. Multiple participants said that raising ethical issues was often seen as disruptive or annoying, for instance, as expressed by this participant: *“I’m already the party pooper, that’s how I feel. Often, when I raise concerns, it is like: ‘Okay, it’s great that you’re questioning how big tech handles your data, good that you mention it, but we really want to move forward with this.’ It’s seen as less relevant.”* (P15). Another participant explains that ethics is often seen as slowing down innovation: *“What you sometimes see is that ethics, doing things responsibly, is perceived as slowing down innovation. ‘Can’t we just do this?’ I see that argument more often from men than from women.”* (P18).

Participants also indicated that there were tensions between ethical and commercial interests of the company. One participant worked on identifying how an algorithm could harm minority groups. She explained what happened when she shared her findings with colleagues who were responsible for technical validation: *“At one point they said: ‘We just don’t have time for your input, we have a budget. We’ve got three or four months, and this is exactly what the client asked for.’ [...] And I thought: well, that kind of undermines the whole point.”* (P16). Some participants mentioned explicitly that AI ethics was not taken seriously, or not even considered ‘real’ AI work, as illustrated by this quote: *“I was at a webinar with top-level women, and they all said: ‘I don’t work in AI.’ At one point I said into the mic: ‘Ladies, I don’t know, but if you don’t work in AI, then who does?’ [...] The reason was that they were focused on responsible AI, not hardcore coding. [...]. Apparently, everything around [hardcore coding] is seen as less important and not even called AI. But this is just as much AI.”* (P18).

The undervaluation of ethical issues was in many cases also reflected in the governance of AI ethics in the companies the participants worked for. Though some of the participants said that their companies had a proper AI ethics governance, which manifested, e.g., in assessments and checks of ethical risks, compliance with legislation, whether records were kept appropriately, whether multiple perspectives were considered, whether data quality was ensured, and whether bias was minimized, many of the participants reported that their companies had no proper AI ethics governance. In part of the companies, often startups, there was no policy regarding AI ethics at all: *“At the moment, no, I regret to say. I think it’s due to the startup culture, very much ‘let’s just do it.’”* (P11). Other companies had AI ethics policies, but they were not actively applied, as illustrated by the following quote: *“I’m sure it’s written down somewhere, but it’s just a general vision document. You can say ‘it must*

be responsible, but that's very vague." (P1). In companies lacking formal policies, some teams or departments had developed their own guidelines or practices. This sometimes led to duplicated efforts, with different departments working independently on AI ethics.

To sum it up, though ethical issues were common in the participants' AI-related work, bringing them up was often not appreciated, as attention for ethics was often seen as hampering innovation and commercial interests. As a result, work on AI ethics was often undervalued and speaking up about ethical issues could be emotionally taxing.

4.3 Women Are Often Taken Less Seriously in AI Workplaces

Although many participants initially said they did not see working as a woman in a male-dominated environment as a problem, throughout the interviews, nearly all were able to share examples of unpleasant experiences or struggles related to their gender. A striking number of participants mentioned being underestimated as a problem they had encountered. This manifested in several ways. One common experience of the participants was being underestimated during first encounters, as illustrated by this quote: *"It's more about the first impression I give. People often think I'm the secretary, the intern, or [...] the coffee lady. I've also given training sessions, and the men I was training, who were usually older, genuinely thought I was there to bring coffee in the morning."* (P14). People often assumed that women held lower-ranking or less technical roles than they actually did, e.g.: *"When people see our team of five, one product manager and four data scientists, I'm always asked: 'Oh, you're not the product manager?' No, that's my colleague. [...] Those kinds of assumptions [...] are a bit irritating."* (P1).

The underestimation of the participants did not stop after first impressions. Many of them said they had to work harder than their male colleagues to be taken seriously. This took various forms. Participants told about how their input was dismissed or ignored: *"I've had several instances where I asked a male colleague something and didn't get a response, but when a male colleague asked the same thing, he did. There was also a time when I had an error, and so did a male colleague. They told me it must be my fault, but not him, and it turned out to be the same issue."* (P12). They told their intelligence was underestimated: *"My intelligence is often underestimated. [...] When a man says the same thing, suddenly people listen."* (P10). They gave examples of situations in which they were given less space to speak: *"I once gave a presentation at a dinner event. At the table afterward, someone just wouldn't take me seriously. They'd talk over me and had no interest in hearing what I had to say."* (P9). They gave examples of their competence being questioned because of their gender: *"Potential clients have said things like: 'Aren't there more men at that company? If*

it's mostly women, will they be able to deliver?'" (P8). They talked about being subjected to belittling remarks: *"It often comes down to men making derogatory comments about women, like saying they're not technical. That makes me uncomfortable, and it's not easy to push back when you're one of the few women."* (P11). And finally, they gave examples of being outright ignored: *"At a client in Nigeria, I wasn't addressed at all, even though I led the project. They directed all their questions to my male colleagues."* (P6) and *"The worst is when [name of male colleague] and I go somewhere together. We were at a diversity and inclusion event, and a man came up to us. [...] He only asked [name of male colleague] questions about the company and what we do. I'm the only tech expert and the only one who knows about the topic, but he just ignored me, like, 'She's a woman, she must be from marketing or HR or something.'"* (P8).

Next to explicitly being underestimated, not taken seriously or ignored, most of the participants had encountered inappropriate comments about their appearance or sexually suggestive remarks at work, which is a more implicit form of undermining being taken seriously as a professional. These are two examples of comments to participants because of their gender, undermining their professionalism: *"I once received a comment during a training I was giving. A group of men from the same company said something sexually suggestive like: 'Oh, you're wearing a skirt today, we were hoping you would, we talked about it yesterday.' [...] That made me feel uncomfortable."* (P14) and *"In the past, I got a lot of comments about my appearance, clothing, things like: 'If I were 20 years younger, I'd know what to do.' I was 25, he was 60. That's just not okay."* (P10). Other participants described intrusive behavior from male colleagues or clients, e.g.: *"Unpleasant situations where younger women were placed at client sites, mostly male IT departments, and the men didn't handle it well. They'd send inappropriate WhatsApp messages, clearly crossing the line. Things like hearts and kisses or repeatedly asking them out for drinks after work. That's just not okay."* (P5). In these examples, women's professional identities were overshadowed by gendered social dynamics, which undermined them to be taken seriously as a professional.

The quotes shown in this paper are a small sample of examples of inappropriate behavior the participants provided. Other examples included awkward jokes, inappropriate comments about menstruation, insensitive remarks about pregnancy, comments on clothing or appearance, assumptions about how women would be perceived by clients, and expectations that women would handle social tasks like organizing team events or bringing cake. There were also cases of men explaining things to women who were equally or more knowledgeable, and women being paid less than male colleagues for the same work.

In nearly every interview, participants mentioned other identity factors than gender that influenced how they were

treated. The following participant reflected on the combination of gender and age: *"I have to work harder to get the same response. But I find it hard to say whether that's purely about gender, because there's also an age difference between me and some of my male colleagues. The youngest is 30, and some are 40 or 50. Maybe they just don't see me as senior enough."* (P12). Many of the participants in more junior roles thought they were taken less seriously because of their age. One of the participants who held a more senior role said that her age helped to be taken more seriously now than in the past: *"In recent years, I've noticed that age has a positive effect. Because I have experience, I'm taken more seriously than someone who says the same thing at 22."* (P7). Another participant spoke about how the combination of gender and religious identity undermined her recognition, specifically, her wearing a headscarf: *Sometimes I think it works in my favor, being a woman and also wearing a headscarf. [...] But other times I think people don't realize I have a technical background and that I'm just as technically skilled and capable."* (P16). A third participant discussed how the intersection of gender, age, and skin color made it more difficult for her to be taken seriously: *"I think that somehow, even though people don't explicitly think about it when they see me, they have this implicit or deeply ingrained feeling that I'm somehow lesser. [...] I have less power. I think it's because I'm not white, not male, and not older. I'm the opposite of all those things."* (P17). Next to age, religion and ethnicity, participants mentioned the following other identity factors influencing how they were treated: cultural background, extraversion level, educational level, educational background (technical vs. non-technical), parents' educational level, parents' income, physical ability, and sexual orientation.

Several participants noted that not only women's identities mattered in how seriously they were taken, but also those of male colleagues. Age was mentioned several times in this context, participants saying that they experienced less sexism from younger men than from older ones: *"It's changing. I especially notice it with the younger generation. The men, 30 and under, they're different. They value my opinion just as much."* (P10).

Unpleasant comments and behaviors have a negative impact on those affected. Several participants said they vividly remembered uncomfortable experiences, even long after they occurred. The following participant reflects on how she responded to a comment she considered inappropriate: *"Now I think: yeah, that really wasn't okay, I should've said something. [...] But I still remember it very clearly, especially because I regret not doing anything about it."* (P10). Even though the participant was not responsible for the inappropriate comment, it continued to occupy her thoughts. Her attention focused not only on the comment itself but also on her own reaction. This illustrates how such experi-

ences can consume mental energy; energy that could otherwise be spent on the substantive aspects of the job. Many participants said they were constantly aware of being in the minority. The following participant gives an example of what that can look like: *"Whenever I walk into a room, I always, consciously or unconsciously, count how many women there are."* (P5). This example shows that simply being part of a minority group can consume mental capacity, even if one does not experience inappropriate behavior.

Negative experiences in the workplace can have long-term effects. One participant described how these experiences subtly undermined her confidence over time: *"It's all very implicit, especially in the earlier stages of my career in tech and IT environments. Comments about clothing, appearance, whether you'd come across well to clients, whether clients would take you seriously, your career prospects. It's always these little pinpricks that slowly, subconsciously, but deeply, make you insecure."* (P5).

In summary, the interviews strongly suggest that being underestimated or not taken seriously is a widespread issue for women working in the Dutch AI sector. Almost all participants had experienced it in one form or another. The results also suggest that being taken seriously as a woman intersects with other identity characteristics, where deviating from the norm makes it more difficult to be taken seriously. Finally, the results suggest that not being taken seriously can cause harm in the short term as well as in the long term, and that it can seriously undermine women in their work.

4.4 Homogeneity Is Reinforced in Large Parts of the AI Sector

Many of the participants described working or having worked at a company with a male-dominated workforce and a non-inclusive or masculine working culture. Based on the participants' responses, we identified three ways in which homogeneity in workforce composition as well as in work culture was reinforced in many of these companies.

First, one way to counter homogeneity in workforce composition and work culture is by installing and complying with diversity and inclusion policies. However, only some of the participants said their companies had successful diversity and inclusion policies. Other participants told that their companies had little to no policy in place, especially participants who worked at smaller companies, e.g.: *"There are plenty of other problems, like startup culture. [...] No, we've never talked about gender ratios or other gender issues."* (P11). Most participants said that their company did have diversity and inclusion policies, but that these were ineffective. One participant said that her company had introduced policies but lacked the capacity to implement them properly. Another described an internal study on the career progression of women, which revealed pay disparities in senior roles. Three years later, however, nothing had been

done with the findings. Some interventions were seen as misguided, for example, those targeting only women: *“A course on female leadership or something like that, it always feels like: ‘Let’s give women tips on how to better navigate this male environment.’”* (P15). Several participants criticized this gap between policy and practice, e.g.: *“The worst thing you can do is have a policy, say all the right things, and then not follow through. That’s basically a license to do the opposite.”* (P18). Thus, homogeneity in the AI sector seems to be kept intact due to a lack of successfully implemented diversity and inclusion policies in a large part of the Dutch AI industry.

Second, homogeneity in the AI sector also seems to be reinforced by how women respond to masculine working environments, i.e., by adapting their behavior to the dominant norms in the working culture. For example, several participants said they adjusted their behavior or appearance to be taken more seriously or to fit in better, e.g.: *“I always make sure I look very neutral. If I wear something colorful, I’m taken less seriously. [...] I even lower my voice.”* (P3). Another participant explained how she adjusted her behavior when working in IT workplaces: *“I worked in IT for a long time, more on the technical side of the company. That’s a very masculine culture, and I learned that when you’re new, you should always crack a joke. That way, you’re immediately one of the guys.”* (P18). Some of the participants initially said they had not experienced any discrimination or inappropriate behavior in their work but started to recall examples throughout the interviews. There were also cases in which participants did not perceive sexist behavior as too problematic: *“Some clients, especially older men, sometimes say things like: ‘I’m glad [name] is here, it gives me something nice to look at.’ I don’t find that too bothersome myself, but I can imagine other women wouldn’t like it.”* (P14) and *“I haven’t been heavily discriminated against, but it’s the kind of stuff we just accept. Like pay gaps, talents not being recognized, or communication styles that are accepted from men but not from women.”* (P18). These last two quotes reflect a pattern of normalization, in which inappropriate behavior or discrimination becomes so common that even women begin to see it as normal, which is another form of adaptation.

Some participants also mentioned that they adapted their behavior to the existing company culture regarding AI ethics. Several participants mentioned that they purposely did not speak out about ethical issues. Participant 4, for instance, did not feel comfortable ethical raising issues with management or clients: *“Yes, at work, but I’ve never discussed [ethical concerns] with clients. Also, because I think it wouldn’t be appreciated.”* (P4). Another participant admitted that although she sometimes felt ethical discomfort, she enjoyed her work so much that she chose to set those concerns aside. Yet another participant mentioned that she struggled with working for certain clients, such as the clients in the defense

industry, and with how clients used the AI applications she and her colleagues had developed, which was unethical in her view. Still this participant chose to continue working under these circumstances.

Third, homogeneity in workforce composition and working culture also seems to be reinforced because of women’s career choices in response to masculine work environments. Part of the participants reacted to negative experiences in their workplace by leaving it. The following quote describes such a situation: *“A colleague spoke up about not being given the floor in meetings by her male teammates. She was in an all-male team, and she’s a woman with dark skin. The team reacted very negatively and aggressively. She ended up switching teams because the manager didn’t do anything about it.”* (P12). Several participants in the study had left unpleasant working environments and found better alternatives. This is illustrated by the following participant, who said that she considered company culture when choosing a job: *“One of the reasons I chose [company] is because it’s extremely diverse, and that really excites me. I think the male-female ratio in this area is quite good, maybe even 50/50, which is awesome.”* (P19). A participant that worked as a freelancer said that she selected clients based on company culture: *“Most of the organizations I work for now, and I choose them deliberately, take this [culture] seriously, because it aligns with my own values.”* (P13).

Some of the participants in more senior roles had been able to create workplaces in which they could strongly influence workforce composition and work culture. One of the participants who was the director of a large department stated: *“At the time, the department had about 25 people. Including management, there wasn’t a single woman. I found that really strange. So I immediately made changes. [...] Next year, I’m planning to roll out an awareness training for everyone, not just my leadership team, but the whole community, to help everyone level up on inclusion.”* (P5). Another participant had founded her own company, in which she created a more diverse and inclusive work culture: *“That’s what really motivates me not to leave tech, even though I considered it for a moment. I want to build something where that kind of exclusion doesn’t happen, where everyone feels welcome.”* (P8). Like this last participant, many others said that they felt a personal responsibility to contribute to gender equality in the field. Several saw themselves as role models and actively advocated for more diversity and inclusion in the sector.

In conclusion, the interview results indicate that there is a substantial number of companies that do not seriously address the widespread issue of women taken less seriously in the AI sector through diversity and inclusion policies or otherwise. Many of the women working in non-inclusive environments in this study responded to that by adapting to masculine norms, making the culture more homogenous, or by

leaving them, which also reinforces homogeneity. These dynamics are likely to contribute to the AI industry remaining homogeneous in terms of both the composition of its workforce and its work culture.

5 Discussion

In this study, we conducted interviews with 19 women who worked in the Dutch AI industry to better understand the role of gender in AI ethics in practice. An important limitation of this study is that it is based on interviews with 19 women, which is a relatively small number. On top of that, selection bias may be present. All participants responded to a recruitment call for research on the role of gender in AI ethics, which could mean they were more interested in gender issues or AI ethics than the average female AI worker. Some participants indeed expressed a strong interest in these topics, but others said they enjoyed working in male-dominated environments and disliked framing their work in terms of gender differences. Another form of possible selection bias is that all participants were working in the AI sector at the time of the interview. Though the recruitment call included people who *had* worked on or with AI, there were no participants who had left the sector in this study, even though research shows that women are more likely to exit the tech field than men (Proebster and Marsden 2026). This could mean that the study paints a more positive picture of the AI working culture for women than reality, as women with more negative experiences may already have left the field. Moreover, multiple participants in the study had actively worked on achieving a more inclusive working environment, e.g., by changing jobs, becoming a freelancer able to select their own clients, or by starting their own company in which they could create their desired company culture. A final limitation of this study is that we only included women and no other genders, while particularly the perspectives of non-binary or gender-fluid people's perspectives could be very valuable (Alfrey and Twine 2017).

Despite these limitations, the participants in our study varied across multiple dimensions, as they worked in different types of companies (freelancers, SMEs, large corporations), represented various sectors (e.g., consultancy, finance, logistics), held diverse roles (e.g., programmer, data analyst, manager), had varying levels of seniority (junior to managing director), and displayed cultural diversity. This suggests that the participant group represented a broad cross-section of women working with or on AI, painting a rich picture of the AI sector in The Netherlands.

A first important outcome of this study concerns the first of the four key themes that we defined, i.e., women enrich AI ethics. Within this theme, we identified three ways brought up by participants in which they believed women

enrich AI ethics work: women (and people from other minority groups) are more likely to recognize ethical issues related to AI, women more often work on AI ethics, and women add additional perspectives to ethical issues related to AI. Though these findings come from a study with a relatively small number of participants, they are in line with similar findings in earlier work. For example, existing work has shown that female AI-scholars more often work on the ethical and societal impacts of AI than their male colleagues (Stathouloupoulos and Mateos-Garcia 2019). As mentioned in the background section, it has already been argued that those who are more likely to be harmed by AI systems, are more likely to recognize ethical risks in these systems and more willing to put effort in mitigating the risks (West et al. 2019; Zowghi and da Rimini 2023), and that women tend to address a broader range of ethical AI issues than men (Hagendorff 2020). However, these studies were not based on empirical research with women working in the AI industry. As such, this study strengthens the existing body of work on how women enrich AI ethics.

A second noteworthy outcome of this study emerges when the first three key themes are considered in relation to each other. The second theme is that ethics is undervalued in the AI sector. This is not a new observation, as multiple authors have argued that current AI ethics practices are often toothless and that there is a lot of so-called ethics-washing in the field (Munn 2023; Rességuier and Rodrigues 2020). The third theme defined in this study is that women are often taken less seriously in AI workplaces. Again, this observation is not new, as it aligns with existing work showing that AI work environments are often not inclusive for women (Shams 2025), and a large body of research on how women in computing, programming and technology are disadvantaged due to harmful stereotypes (Camp 2012; Hill et al. 2010; Margolis and Fisher 2002). Yet, when connecting these two themes to the first theme, women enriching AI ethics, a less often highlighted picture emerges. According to this study, the people in the AI industry that are most engaged in AI ethics often seem to be the ones that are least heard. For women speaking up on ethical issues in the AI industry this suggests that they have a double disadvantage: they not only have more difficulty being taken seriously but they also work on a topic that is undervalued, both of which make it more difficult for them to get their message across.

The lower valuation of ethical compared to technical AI work on the one hand, and the lower status of women compared to men in the AI work field on the other hand, seem deeply related. In our study, to illustrate that they were being underestimated and not taken seriously, many participants gave examples of how their *technical* skills and knowledge were underestimated, and how people assumed that they held *less technical* roles than they did. This relates to work by Hermans and Schlesinger (2024), who argue that technical work in computer science, e.g. work that focuses on

features of programming languages, is considered more difficult and has a higher status in the field than less technical work, e.g., work that focuses on applications, adoption or the needs of users. It is important to note that, although technical work is seen as more difficult than work on ethical or societal implications of AI, this is not necessarily the case. For example, one could argue that whereas well-defined technical problems often have clear-cut solutions, problems considering a socio-technical context are seldomly clearly delineated or fully solvable, and therefore, are more difficult to work on. Yet, in practice, the dominant view in computer science is that technical work is more difficult, and therefore, valued higher (Hermans and Schlesinger 2024).

A third important outcome we would like to highlight is an insight that emerges when the fourth key theme, i.e., homogeneity is reinforced in large parts of the AI sector, is considered in light of the interplay between the first three themes, i.e., a seemingly double disadvantage for women speaking up on ethical issues in the AI industry. Our findings related to the fourth theme are in line with existing literature, which shows that AI work environments are often not inclusive for women (Shams 2025), women navigate workplace hostilities by distancing themselves from conventional heterosexual femininity (Alfrey and Twine 2017), and women are more likely to exit the tech field than men (Proebster and Marsden 2026). However, whereas existing work often attributes women's departure from the field to female-unfriendly work environments, our study suggests that, in addition to that, women may also leave the sector because they are interested in working on AI ethics. It has been argued before that, due to the undervaluation of AI ethics, raising ethical issues or working on AI ethics requires extra (emotional) efforts on top of the substantive work (Gambelin 2021). This applies especially to *women* working on AI ethics, because of the double disadvantage. As a result, women who care about AI ethics and put effort in the ethical development and deployment of AI-systems, may get disappointed when their contributions are not heard or taken seriously. They may therefore either stop putting effort in AI ethics or leave their department or company altogether. Thus, the undervaluation of ethics could, in addition to non-inclusive work environments, contribute to the reinforcement of homogeneity in AI workplaces, which in turn, may cause ethics to remain undervalued.

The findings of our study have different implications. For the AI sector, deprioritizing or ignoring ethical concerns increases the risk of scandals and reputational damage. Deprioritizing diversity and inclusion may cause women and other minorities to leave the workplace. This would be a loss, as it worsens ethical decision making (DeGrassi et al. 2012) and, in addition to that, it hinders innovation in general, as research shows that diversity in teams results in more creativity and innovation (Hundschell et al. 2022; Wang et al. 2019), and higher code quality in AI systems (Cynthia

and Roy 2025). Moreover, valuable talent may be lost at a time in which AI workers are in high demand (Leopold et al. 2025). Finally, the homogeneity in workforce composition and work culture may restrict developments in the field, as potentially interesting topics and ideas are not always acknowledged in masculine work cultures (Adam 1998; Hermans and Schlesinger 2024). For society, the findings of this study, indicating dynamics in the AI sector due to which AI ethics remains persistently undervalued, mean that AI systems may keep causing ethical harms and perpetuate discrimination and exclusion through AI systems.

6 Conclusion

This paper presented a study on the role of gender in the ethical development and deployment of AI, based on interviews with 19 women working in the Dutch AI industry. Following a thematic analysis of the interviews, we defined four key themes: women enrich AI ethics; AI ethics is undervalued in the AI sector; women are often taken less seriously in AI workplaces; and homogeneity of workforce composition and work culture is reinforced in large parts of the AI sector. These themes are related to each other; those people that tend to work on AI ethics more, seem to be the ones that are taken less seriously, and as a result, more likely to leave their workplace. It seems that, therefore, non-inclusive work cultures and the undervaluation of ethics can remain intact in large parts of the AI industry.

This study suggests that diversity and inclusion on the one hand, and AI ethics on the other hand, are deeply interconnected. Yet in practice, they are rarely addressed together, meaning that important insights could be missed. Future research could further study how diversity and inclusion, and AI ethics interact with each other in different contexts. In this study we focused on AI industry in The Netherlands. Future research could target other countries, focus on other contexts, e.g. the public sector or academia, or dive deeper into specific types of contexts, e.g. startup culture. Furthermore, our study focused on gender, while there are other identity characteristics that influence a person's experiences in the AI sector, like ethnicity, race, class, ability, age and sexual orientation (Ulnicane 2024). Future work could foreground (the combination of) other identity characteristics and their relation to AI ethics, following an intersectional approach (Crenshaw 1991). A final direction for future work is to study how the current situation can be changed. In our study, there were several participants who worked at or had been able to create AI workplaces that were more inclusive, diverse and in which AI ethics was highly valued. Such places can be pivotal in showing how (gender) equality and AI ethics can be placed into the center of AI development and deployment, creating AI applications that are beneficial not only for a few but for everyone.

Acknowledgments

We would like to thank the team at VHTO (Dutch expertise center for girls and women in STEM) for their valuable feedback and insightful comments. This work was supported by the Province of South-Holland in The Netherlands and by the Dutch Taskforce for Applied Research SIA, under grant number SPR.ALG.01.02.007.

References

- Adam, A. 1998. *Artificial Knowing: Gender and the Thinking Machine*. London: Routledge. doi.org/10.4324/9780203005057.
- Alfrey, L.; and Twine, F. W. 2017. Gender-fluid geek girls: Negotiating inequality regimes in the tech industry. *Gender and Society* 31(1): 28-50. doi.org/10.1177/0891243216680590.
- Benjamin, R. 2019. *Race after Technology: Abolitionist Tools for the New Jim Code*. Oxford: Polity.
- Braun, V.; and Clarke, V. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3(2): 77–101. doi.org/10.1191/1478088706qp063oa.
- Braun, V.; and Clarke, V. 2019. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health* 11(4): 589–597. doi.org/10.1080/2159676X.2019.1628806.
- Braun, V.; and Clarke, V. 2023. Toward good practice in thematic analysis: Avoiding common problems and becoming a knowing researcher. *International Journal of Transgender Health* 24(1): 1–6. doi.org/10.1080/26895269.2022.2129597.
- Buolamwini, J.; and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of Machine Learning Research, Volume 81: Conference on Fairness, Accountability and Transparency*. New York: ML Research Press.
- Camp, T. (2012). 'Computing, we have a problem...'. *ACM Inroads*, 3(4): 34-40. doi.org/10.1145/2381083.2381097
- Crenshaw, K. 1991. Mapping the Margins: Intersectionality, Identity Politics, and Violence against Women of Color. *Stanford Law Review* 43(6): 1241–1299. doi.org/10.2307/1229039.
- Criado Perez, C. 2019. *Invisible Women: Exposing Data Bias in a World Designed for Men*. New York: Random House.
- Cynthia, S. T.; and Roy, B. 2025. An empirical study on the impact of gender diversity on code quality in AI systems. In *Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering*. New York: Association for Computing Machinery. doi.org/10.1145/3696630.3731674.
- DeGrassi, S. W.; Morgan, W. B.; Walker, S. S.; Wang, Y. I.; and Sabat, I. 2012. Ethical decision-making: Group diversity holds the key. *Journal of Leadership, Accountability and Ethics* 9(6): 51-65.
- Ferrer, X.; Van Nuenen, T.; Such, J. M.; Côté, M.; and Criado, N. 2021. Bias and discrimination in AI: a cross-disciplinary perspective. *IEEE Technology and Society Magazine* 40(2): 72-80.
- Gambelin, O. 2021. Brave: what it means to be an AI Ethicist. *AI and Ethics* 1(1): 87-91.
- Gibson, C.; Olszewski, D.; Brigham, N. G.; Crowder, A.; Butler, K. R.; Traynor, P.; Redmiles, E. M.; and Kohno, T. 2025. Analyzing the AI nudification application ecosystem. In *Proceedings of the 34th USENIX Conference on Security Symposium (SEC '25)*. USA: USENIX Association.
- Green, A.; and L. Lamby 2023. The supply, demand and characteristics of the AI workforce across OECD countries. *OECD Social, Employment and Migration Working Papers*, No. 287, Paris: OECD Publishing. doi.org/10.1787/bb17314a-en.
- Hagendorff, T. 2020. The ethics of AI ethics: An evaluation of guidelines. *Minds and machines* 30(1): 99-120.
- Harbers, M.; Horsman, S.; Wiggers, P.; Aldewereld, H.; Fuckner, M.; Van Gemeren, C.; Hajri, O.; Schukking, F.; and Stembert, N. 2025. Responsible AI in Media Organizations: Four Case Studies Implementing Ethical Tools in Practice. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 8(2): 1219–1230. doi.org/10.1609/aies.v8i2.36624.
- Hermans, F.; and Schlesinger, A. 2024. A case for feminism in programming language design. In *Proceedings of the 2024 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. New York: Association for Computing Machinery. doi.org/10.1145/3689492.3689809.
- Hill, C.; Corbett, C.; and St Rose, A. 2010. Why so few? Women in science, technology, engineering, and mathematics. Washington: American Association of University Women.
- Hundschell, A.; Razinskas, S.; Backmann, J.; and Hoegl, M. 2022. The effects of diversity on creativity: A literature review and synthesis. *Applied Psychology* 71(4): 1598-1634.
- Hunkenschroer, A. L.; and Luetge, C. 2022. Ethics of AI-enabled recruiting and selection: A review and research agenda. *Journal of Business Ethics* 178(4): 977-1007.
- Jobin, A.; Ienca, M.; and Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature machine intelligence* 1(9): 389-399.
- Khan, A. A.; Badshah, S.; Liang, P.; Waseem, M.; Khan, B.; Ahmad, A.; Fahmideh, M. Niazi, M.; and Akbar, M. A. 2022. Ethics of AI: A systematic literature review of principles and challenges. In *Proceedings of the 26th International Conference on Evaluation and Assessment in Software Engineering (EASE '22)*. New York: Association for Computing Machinery. doi.org/10.1145/3530019.3531329.
- Kronqvist, A.; and Rousi, R. A. 2022. A quick review of ethics, design thinking, gender, and AI development. *International Journal of Design Creativity and Innovation* 11(1): 62-79. doi.org/10.1080/21650349.2022.2136762.
- Kronqvist, A.; and Rousi, R. 2025. Experience of Gender Among Professionals in Finnish IT Companies. In *Proceedings of the International Conference on Software Business*. Springer. doi.org/10.1007/978-3-031-85849-9_17.
- Leopold, T.; Di Battista, A.; Jativa, X.; Sharma, S.; Li, R. and Grayling, S. 2025. *Future of jobs report 2025*. World Economic Forum.
- Margolis, J.; and Fisher, A. 2002. *Unlocking the Clubhouse: Women in Computing*. Cambridge: MIT Press.
- Morley, J.; Floridi, L.; Kinsey, L.; and Elhalal, A. 2021. From What to How: An Initial Review of Publicly Available AI Ethics Tools, Methods and Research to Translate Principles into Practices. *Science and Engineering Ethics* 26: 2141–2168. doi.org/10.1007/s11948-019-00165-5.
- Munn, L. 2023. The uselessness of AI ethics. *AI and Ethics* 3(3): 869-877.
- Nadeem, A., Abedin, B., & Marjanovic, O. (2020). Gender bias in AI: A review of contributing factors and mitigating strategies. In *Proceedings of CIS 2020*.

- Nadeem, A.; Marjanovic, O.; and Abedin, B. 2022. Gender bias in AI-based decision-making systems: a systematic literature review. *Australasian Journal of Information Systems* 26. doi.org/10.3127/ajis.v26i0.3835.
- Noble, S. U. 2018. *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York university press.
- O'Connor, S.; and Liu, H. 2024. Gender bias perpetuation and mitigation in AI technologies: challenges and opportunities. *AI & SOCIETY* 39(4): 2045-2057.
- Platform Talent voor Technologie. M/V Kenmerken technologische arbeidsmarkt. <https://www.ptvt.nl/dashboard/pagina/my-kenmerken-technische-arbeidsmarkt#grafiek387>. Accessed: 2026-05-19.
- Portegies, T.; Hajri, O.; Willemsen, L.; and Harbers, M. 2025). Decoding AI ethics: What do ethical tools tell us about ethics?. *Ethics and Information Technology* 27(4): 1-12. doi.org/10.1007/s10676-025-09860-3.
- Prem, E. 2023. From ethical AI frameworks to tools: a review of approaches. *AI and Ethics* 3(3): 699-716.
- Proebster, M.; and Marsden, N. 2026. Retention of Women in Tech—A Systematic Literature Review. *Journal of Computer Information Systems*, 1-15. doi.org/10.1080/08874417.2026.2616426.
- Ray, P. P. 2023. ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope. *Internet of Things and Cyber-Physical Systems* 3(1): 121–154.
- Rességuier, A.; and Rodrigues, R. 2020. AI ethics should not remain toothless! A call to bring back the teeth of ethics. *Big Data & Society* 7(2). doi.org/10.1177/2053951720942541.
- Roopaei, M.; Horst, J.; Klaas, E.; Foster, G.; Salmon-Stephens, T. J.; and Grunow, J. 2021. Women in AI: Barriers and solutions. In *Proceedings of the 2021 IEEE World AI IoT Congress (AIIoT)*. IEEE.
- Sanderson, C.; Douglas, D.M.; Lu, Q.; Schleiger, E.; Whittle, J.; Lacey, J.; Newnham, G.J.; Hajkowicz, S.; Robinson, C.J.; and Hansen, D.M. 2021. AI Ethics Principles in Practice: Perspectives of Designers and Developers. *IEEE Transactions on Technology and Society* 4: 171-187. doi.org/10.1109/TTS.2023.3257303.
- Shams, R. A.; Zowghi, D.; and Bano, M. 2025. AI and the quest for diversity and inclusion: a systematic literature review. *AI and Ethics* 5(1): 411-438.
- Stahl, B. C. 2021. Ethical issues of AI. In *Artificial Intelligence for a better future: An ecosystem perspective on the ethics of AI and emerging digital technologies*, 35-53. Cham: Springer International Publishing.
- Stathouloupoulos, K.; and Mateos-Garcia, J. C. 2019. *Gender diversity in AI research*. Available at SSRN 3428240.
- West, S. M.; Whittaker, M.; and Crawford, K. 2019. *Discriminating systems*. AI Now.
- Strengers, Y.; and Kennedy, J. 2021. *The Smart Wife: Why Siri, Alexa, and Other Smart Home Devices Need a Feminist Reboot*. MIT Press. doi.org/10.7551/mitpress/12482.001.0001.
- Trotta, A.; Ziosi, M.; and Lomonaco, V. 2023. The future of ethics in AI: challenges and opportunities. *AI & Society* 38(2): 439-441.
- Ulnicane, I. 2024. Intersectionality in Artificial Intelligence: Framing Concerns and Recommendations for Action. *Social Inclusion* 12: article 7543. doi.org/10.17645/si.7543.
- Wang, J.; Cheng, G. H. L.; Chen, T.; and Leung, K. 2019. Team creativity/innovation in culturally diverse teams: A meta-analysis. *Journal of Organizational Behavior* 40(6): 693-708.
- Young, E.; Wajcman, J.; and Sprejer, L. 2023. Mind the gender gap: Inequalities in the emergent professions of artificial intelligence (AI) and data science. *New Technology, Work and Employment* 38(3): 391-414.
- Zhou, J.; and Chen, F. 2023. AI ethics: from principles to practice. *AI & Society* 38(6): 2693-2703.
- Zowghi, D.; da Rimini, F. 2023. Diversity and inclusion in artificial intelligence. arXiv:2305.12728.