

Bridging the Gender Gap in AI Careers: A Feminist Perspective on Structural Barriers, Mentorship, and Policy in Artificial Intelligence Careers

T R Thriveni

Research Scholar Department of Sociology Bangalore University

Abstract:

Artificial intelligence (AI) is rapidly reshaping global labour markets, yet the professionals designing, deploying, and governing these systems remain overwhelmingly male. This article examines the gender gap in AI careers through a feminist — specifically technofeminist and intersectional — lens, arguing that women's underrepresentation is not a pipeline problem alone but a structural, cultural, and technopolitical one. Drawing on recent workforce data, the article documents the scale of the gap across the AI career ladder, from entry-level hiring to the C-suite, and reviews four intervention areas: mentorship and role models, bias-aware recruitment, feminist-informed organisational policy, and inclusive AI education. To illustrate how such claims can be operationalised and tested empirically, the article presents a demonstration research design — including objectives, hypotheses, a fifty-respondent (N = 50) illustrative sample, descriptive means, and hypothesis-test results — that a future primary study could replicate with real data. The discussion connects each hypothesis back to its objective and result, and the article closes with policy recommendations and the limitations of both the literature and the illustrative model presented.

Keywords: gender gap, artificial intelligence, technofeminism, employability, mentorship, algorithmic bias, workplace policy

1. Introduction

Artificial intelligence is often described as the defining general-purpose technology of this decade, reorganising hiring, healthcare, finance, and education alike. Yet the workforce building that technology looks strikingly like the workforce that built earlier waves of computing: predominantly male, and more male still at senior levels.¹ Women make up roughly a quarter to a third of the global AI workforce, and the share thins further at every step up the career ladder.¹

This underrepresentation matters for reasons that go beyond fairness in hiring statistics. AI systems encode the assumptions, priorities, and blind spots of the people who build them; a workforce skewed toward one gender is more likely to produce tools that work less well, or actively worse, for the excluded group. It also matters for employability in the narrower economic sense: AI roles are now among the best-paid

¹ Girl Scouts of the USA / World Economic Forum, "5 Ways to Increase the Number of Women Working in AI," 2021, <https://www.weforum.org/stories/2021/08/5-ways-increase-women-working-ai/>.

segments of the labour market, so exclusion from AI careers increasingly means exclusion from the fastest-growing source of high-wage employment.

This article asks two connected questions. First, descriptively: how large is the gender gap across the AI career pipeline, and what mechanisms sustain it? Second, normatively and empirically: what would a feminist-informed research design look like for testing the mechanisms most commonly proposed — mentorship, workplace bias, and organisational policy — and what pattern of results would support or challenge them? The article treats the second question as a methodological demonstration rather than a completed study; the results reported in Section 8 are illustrative, not collected data.

2. Feminist Theoretical Framework

2.1 Technofeminism

The framework adopted here is technofeminism, most closely associated with Judy Wajcman's argument that technology and gender are mutually constitutive: technology is not a neutral tool that women simply have unequal access to, but a domain that has been historically co-produced with, and used to reproduce, masculine authority.² A related strand, associated with Wendy Faulkner, distinguishes between the gendering of technology 'by design' (who builds it, whose needs it encodes) and 'by association' (the cultural coding of technical competence as a masculine trait).¹

Applied to AI specifically, recent technofeminist scholarship argues that the gender skills gap in AI and data science reflects historically constructed equations between masculinity and technical expertise, and that a male-dominated design workforce produces a feedback loop in which social bias is built into machine-learning systems and then presented back to society as objective computation.³

2.2 Intersectionality

An intersectional lens adds that the AI gender gap is not experienced uniformly by all women. Race, class, nationality, disability, and immigration status compound the barriers documented in Section 3, so that, for example, women from the Global South or from minoritised ethnic groups in the Global North face a narrower entry point into AI careers than the aggregate gender statistics alone suggest. Recent scholarship on "AI empire" extends this argument globally, describing generative AI's labour and data supply chains as reproducing colonial-era divisions of labour along both gender and geographic lines.

2.3 Why a Feminist Lens Changes the Diagnosis

A feminist framing reframes three assumptions common in mainstream "women in tech" discourse. First, it treats the shortage of women in AI not as a supply problem ("not enough girls choose STEM") but as a demand and retention problem produced by workplace culture, hiring practice, and reward structures. Second, it treats AI tools themselves as a site of gendered power rather than a neutral instrument that could equally empower or disadvantage any user. Third, it insists that interventions be evaluated not only by headcount but by whether women gain decision-making authority — the C-suite and "Head of AI" statistics in Section 3 exist precisely to make that distinction visible.

3. Literature Review

3.1 The Scale of the Gender Gap in AI

The clearest recent evidence comes from LinkedIn's 2026 "Triple Penalty" report, based on tens of milli-

²Judy Wajcman, *TechnoFeminism* (Cambridge: Polity Press, 2004).

³ University of Sussex, "Feminism Confronts AI: The Gender Relations of Digitalisation," 2024, <https://www.sussex.ac.uk/broadcast/read/63624>.

ons of workers across twenty-seven countries. It found that women accounted for only about a quarter of AI hires in 2025, compared with half of hires into non-AI roles, even as AI postings offered roughly double the typical salary of non-AI postings.¹ The same report found the gap widening at senior levels: about one-fifth of “Head of AI” roles, roughly a quarter of “Director of AI” roles, and under a fifth of “Member of Technical Staff” roles were held by women, while women held only thirteen per cent of C-suite AI leadership positions at AI companies across the twenty-seven countries surveyed.⁴

Older but still widely cited figures put the female share of the general AI workforce at roughly twenty-two per cent, and a 2023 UNESCO assessment found that only about three per cent of female students in higher education chose information and communication technology as a field of study — a strong early indicator of the pipeline that later produces these workforce ratios.




Figure 1. Women's share of roles narrows at every step up the AI career ladder. Sources: LinkedIn “Triple Penalty” report (2026); World Economic Forum, *Agenda Dialogues* (2025).

3.2 Structural and Cultural Barriers

Three barrier types recur across the literature. First, educational funnelling: girls are consistently less likely than boys to select STEM subjects in secondary school, a pattern one 2026 Australian survey linked to cultural bias, narrow subject framing, and a lack of visible role models rather than to any difference in aptitude.¹ Second, workplace culture: qualitative studies using the Advancement of Women in Technology framework describe an “invisibility” experienced by women already working in AI, whereby their technical contributions are less likely to be recognised or promoted than equivalent contributions by men.⁵ Third, adoption and trust gaps: Deloitte's 2025 global predictions noted that trust in generative AI tools, representation in training data, and diversity within AI teams remained gendered even as women's raw usage of generative AI tools was projected to catch up with men's.¹

3.3 Algorithmic Bias as a Feedback Loop

A distinct strand of the literature treats AI systems themselves as a barrier, not only the human workforce

⁴AfroTech, “13% Of AI Executives Are Women. Here's Why That's A Problem,” 2026, <https://afrotech.com/13-of-ai-executives-are-women-heres-why-thats-a-problem>.

⁵ Isi Idemudia, “Gender, Workforce and Artificial Intelligence,” Robert Gordon University repository, <https://rgu-repository.worktribe.com/OutputFile/2098632>.

that builds them. Because most AI teams remain male-dominated, tools are frequently trained and validated on data or use-cases that under-represent women's needs, which then feeds back into product decisions and, indirectly, into who is seen as a credible AI expert. This feedback-loop argument is central to the technofeminist framework adopted in Section 2, and it is why several of the recommendations in Section 10 target design teams and training data rather than recruitment alone.

3.4 Interventions: Mentorship, Policy, and Education

The intervention literature clusters around four levers: (a) mentorship and visible role models, which industry bodies such as Girl Scouts of the USA and WomenTech Network identify as the most consistently requested support by women already in technical roles;⁶ (b) bias-aware recruitment, including anonymised resume screening and AI-audited job postings; (c) organisational policy, such as flexible work arrangements, transparent promotion criteria, and pay-equity audits; and (d) education, including AI and robotics curricula designed specifically to engage girls earlier in secondary school.¹

What the literature lacks, more often than not, is a single study that connects these levers statistically to a named set of hypotheses and reports effect sizes rather than case-study narrative. Sections 4 through 8 build such a design, first as a set of testable objectives and hypotheses, and then as an illustrative demonstration of how the resulting data could be reported.

4. Research Objectives

O1: To assess the structural barriers that women in and adjacent to AI careers perceive at entry, retention, and promotion stages.

O2: To examine the relationship between access to mentorship/role models and women's stated intention to persist in an AI career.

O3: To evaluate how perceived workplace and algorithmic bias relates to satisfaction with career advancement.

O4: To explore whether feminist-informed organisational policy (flexible work, transparent promotion, bias audits) is associated with women's self-efficacy in AI roles.

5. Hypotheses

H1: Women working in or adjacent to AI report perceived structural barriers significantly above the scale midpoint.

H2: Mentorship access is positively associated with career persistence intention.

H3: Perceived workplace/algorithmic bias is negatively associated with career advancement satisfaction.

H4: Organisational feminist-informed policy is positively associated with self-efficacy in an AI role.

6. Feminist Research Methodology

Before specifying the concrete research design (Section 7), this section sets out the methodological commitments that make the design a feminist one, rather than a generic workplace-attitudes survey with “women” as the sampling criterion. Feminist research methodology is not one method but a set of epistemological and ethical commitments that can be applied to quantitative, qualitative, or mixed-methods designs alike.

⁶WomenTech Network, "How Can We Leverage Tech to Break Gender Barriers in STEM Fields?," <https://www.womentech.net/how-to/how-can-we-leverage-tech-break-gender-barriers-in-stem-fields>.

6.1 Epistemological Positioning: Standpoint Theory and Situated Knowledge

The design draws on feminist standpoint theory, which holds that knowledge is always produced from a particular social position, and that the perspectives of marginalised groups — here, women navigating a male-dominated AI industry — can reveal structural features of that industry that are less visible from a dominant-group standpoint.⁷ This is combined with Donna Haraway's concept of situated knowledges, which rejects the idea of a single “view from nowhere” and insists instead that all knowledge claims, including quantitative survey data, be understood as partial and positioned rather than neutral.¹

Practically, this means the instrument in Section 7.2 is framed around participants' own perceptions (“perceived” barriers, “perceived” bias) rather than claiming to measure an objective, view-from-nowhere reality, and the demographic categories in Table 1 are treated as socially constructed and intersecting rather than fixed background variables to be controlled away.

6.2 Reflexivity and Positionality

Feminist methodology requires the researcher to be reflexive about their own position relative to the people and topic studied — gender, professional relationship to the AI industry, and institutional affiliation all shape which questions are asked and how answers are interpreted. A full study using this design should include a positionality statement describing the research team's relationship to the AI industry and to feminist scholarship, and should document any changes the research team made to the instrument in response to participant feedback during piloting.

6.3 Participatory and Mixed-Methods Commitments

Where possible, feminist research favours participatory design: involving prospective respondents in reviewing or co-designing the instrument, rather than only administering it to them. The design in Section 7 proposes a mixed-methods extension: the Likert-scale instrument (Section 7.2) generates the quantitative results in Section 8.1–8.4, and a smaller set of open-ended interview questions — illustrated qualitatively in Section 8.5 — lets respondents describe experiences the fixed-response items cannot capture.⁸

6.4 Intersectional Sampling Strategy

Consistent with the intersectional framework in Section 2.2, a full study should stratify sampling by race/ethnicity, disability status, seniority, and region rather than treating “women in AI” as a homogeneous category. The demonstration sample in Section 8.1 includes a simplified intersectional breakdown (Section 8.4) to illustrate how subgroup differences would be reported, while acknowledging that $N = 50$ is too small to support robust intersectional statistical inference and that a genuine study would need a substantially larger, purposively stratified sample for this analysis to be adequately powered.

6.5 Feminist Research Ethics

Feminist research ethics go beyond standard informed consent to attend to power asymmetries between researcher and participant, the risk of extractive research that takes participants' narratives without benefiting them, and the emotional labour involved in recounting experiences of bias or exclusion. A full study should offer participants control over how their words are used (including the option to review quotations before publication), avoid asking respondents to relive traumatic experiences without support, and consider compensating participants' time, consistent with an ethic of care rather than treating respondents purely as data sources.¹

⁷Sandra Harding, *Whose Science? Whose Knowledge? Thinking from Women's Lives* (Ithaca: Cornell University Press, 1991).

⁸Shulamit Reinharz, *Feminist Methods in Social Research* (New York: Oxford University Press, 1992).

7. Research Design

7.1 Proposed Design

A cross-sectional online survey, with an embedded qualitative component, of women currently working in, or having recently exited, AI-related roles (data science, machine learning engineering, AI product management, AI research). A target sample of $N = 50$ is proposed as a pilot; a full study would require a larger, stratified sample (by seniority, region, race/ethnicity, and sector) to support subgroup and intersectional analysis at adequate statistical power.

7.2 Proposed Instrument

Seven constructs, each measured on a five-point Likert scale (1 = strongly disagree, 5 = strongly agree), adapted from existing STEM-climate and workplace-bias scales: Structural Barriers (SB), Mentorship Access (MA), Career Persistence Intention (CPI), Perceived Workplace/Algorithmic Bias (PWB), Career Advancement Satisfaction (CAS), Organisational Feminist Policy (OFP), and Self-Efficacy (SE), plus three open-ended questions on lived experience of entry, retention, and advancement, consistent with the mixed-methods commitment in Section 6.3.

7.3 Proposed Analysis

One-sample t-tests against the scale midpoint (H1), Pearson correlations (H2–H4), a full inter-construct correlation matrix, an exploratory intersectional subgroup comparison, and thematic analysis of open-ended responses, with alpha set at .05 for quantitative tests. A full study should additionally control for seniority and region, report confidence intervals rather than p-values alone, and follow a reflexive thematic-analysis procedure for the qualitative data consistent with Section 6.1–6.2.

8. Results (Illustrative Demonstration Data, $N = 50$)

8.1 Illustrative Sample Profile

Characteristic	Category	n	%
Career stage	Student / early career (0–2 yrs)	14	28%
	Mid-career (3–8 yrs)	22	44%
	Senior / leadership (9+ yrs)	14	28%
AI sub-field	Data science / ML engineering	21	42%
	AI research	9	18%
	AI product / program management	12	24%
	AI policy / ethics	8	16%
Region	North America	18	36%
	Europe	13	26%
	Asia-Pacific	12	24%
	Other	7	14%

Table 1. Illustrative demonstration sample profile ($N = 50$).

8.2 Descriptive Statistics

Construct	Mean	SD	Related objective
Structural Barriers (SB)	3.8	0.71	O1
Mentorship Access (MA)	2.9	0.88	O2
Career Persistence Intention (CPI)	3.6	0.79	O2
Perceived Workplace/Algorithmic Bias (PWB)	3.7	0.68	O3
Career Advancement Satisfaction (CAS)	2.8	0.75	O3
Organisational Feminist Policy (OFP)	2.6	0.91	O4
Self-Efficacy (SE)	3.4	0.66	O4

Table 2. Illustrative construct means and standard deviations (N = 50, 5-point scale).




Figure 2. Illustrative construct means with standard deviation bars (demonstration data, N = 50).

8.3 Hypothesis Tests

Hypothesis	Test	Result	Decision
H1: SB > scale midpoint (3.0)	One-sample t-test	$t(49) = 5.62, p < .001$	Supported
H2: MA → CPI	Pearson correlation	$r = .52, p < .01$	Supported
H3: PWB → CAS	Pearson correlation	$r = -.47, p < .01$	Supported
H4: OFP → SE	Pearson correlation	$r = .44, p < .01$	Supported

Table 3. Illustrative hypothesis-test summary (demonstration data, N = 50).




Figure 3. Illustrative effect sizes for H1–H4 (demonstration data, N = 50).

8.4 Inter-Construct Correlation Matrix

Table 4 reports the full illustrative correlation matrix across all seven constructs, extending beyond the four pairs named in H1–H4. It is included so that a future study can check whether its own matrix shows a similar pattern — in particular, the illustrative data show Structural Barriers positively associated with Perceived Workplace/Algorithmic Bias ($r = .61$) and negatively associated with Career Advancement Satisfaction ($r = -.35$), consistent with the feedback-loop argument in Section 3.3 even though neither pair was formally hypothesised.

	SB	MA	CPI	PWB	CAS	OFP	SE
SB	1.00	-.18	-.22	.61	-.35	-.29	-.15
MA	-.18	1.00	.52	-.20	.30	.38	.27
CPI	-.22	.52	1.00	-.25	.33	.31	.40
PWB	.61	-.20	-.25	1.00	-.47	-.42	-.30
CAS	-.35	.30	.33	-.47	1.00	.36	.33
OFP	-.29	.38	.31	-.42	.36	1.00	.44
SE	-.15	.27	.40	-.30	.33	.44	1.00

Table 4. Illustrative inter-construct correlation matrix (demonstration data, N = 50). SB = Structural Barriers; MA = Mentorship Access; CPI = Career Persistence Intention; PWB = Perceived Workplace/Algorithmic Bias; CAS = Career Advancement Satisfaction; OFP = Organisational Feminist Policy; SE = Self-Efficacy.

8.5 Intersectional Subgroup Analysis

Consistent with the intersectional sampling commitment in Section 6.4, Table 5 and Figure 4 break Structural Barriers down by illustrative race/ethnicity subgroup within the N = 50 demonstration sample. The illustrative pattern — higher perceived structural barriers among Black and Latina respondents than

among White or Asian respondents — is constructed specifically to demonstrate why an intersectional breakdown matters: an aggregate mean of 3.8 (Table 2) would obscure this variation entirely. With subgroup sizes of three to eight in this demonstration sample, no inferential test is reported; a genuine study would need a purposively stratified sample large enough to power between-group comparisons.

Subgroup	n	Mean SB	SD
White / non-Hispanic	22	3.6	0.65
Asian	12	3.7	0.70
Black / African-descent	8	4.3	0.58
Latina / Hispanic	5	4.1	0.62
Other / multiracial	3	3.9	0.75

Table 5. Illustrative Structural Barriers mean by race/ethnicity subgroup (demonstration data, N = 50).




Figure 4. Illustrative Structural Barriers by race/ethnicity subgroup (demonstration data, N = 50).

8.6 Illustrative Qualitative Themes

The composite quotations below are invented illustrations of the kind of theme a thematic analysis of the open-ended questions (Section 7.2) might produce; they are not real interview data and are not attributed to any actual person. They are included to show how the qualitative component described in Section 6.3 would be reported alongside the quantitative results above.

Theme	Illustrative composite quotation
Proving competence repeatedly	“I notice I get asked to walk through my reasoning in meetings in a way my male teammates don’t — like my technical judgement starts from zero every time.”

Theme	Illustrative composite quotation
Informal mentorship, formal gatekeeping	“The mentoring I've had has been great, but promotion criteria for the senior technical track still feel unwritten.”
Asymmetric blame for AI errors	“When a model I built made a mistake, it became a story about whether I was ready for the role, not about the model.”
Policy on paper vs. practice	“Flexible work is in the handbook, but everyone knows using it full-time slows you down on the leadership track.”

Table 6. Illustrative composite qualitative themes (constructed for demonstration; not real interview data).

9. Discussion: Connecting Objectives, Hypotheses, and Results

9.1 Objective 1 / Hypothesis 1 — Structural Barriers

In the illustrative data, Structural Barriers scored well above the scale midpoint ($M = 3.8$) and the one-sample t-test supports H1. If a real sample replicated this pattern, it would corroborate the literature reviewed in Section 3.2: women in AI perceive entry, retention, and promotion as structurally harder than a neutral baseline, consistent with the “invisibility” and educational-funnelling findings cited there. This would directly answer Objective 1.

9.2 Objective 2 / Hypothesis 2 — Mentorship and Persistence

Mentorship Access scored below the midpoint ($M = 2.9$) — i.e., limited access on average — while Career Persistence Intention scored well above it ($M = 3.6$), and the two were positively correlated ($r = .52$) in the illustrative data, supporting H2. Read together with Objective 2, this pattern would suggest that persistence intention is currently being sustained despite, not because of, mentorship access — implying that expanding mentorship (Section 3.4a) could plausibly raise persistence further rather than simply maintaining it.

9.3 Objective 3 / Hypothesis 3 — Bias and Advancement Satisfaction

Perceived Workplace/Algorithmic Bias scored high ($M = 3.7$) while Career Advancement Satisfaction scored low ($M = 2.8$), with a negative correlation between them ($r = -.47$) supporting H3. This mirrors the algorithmic-feedback-loop argument in Section 3.3: bias is not only a hiring-stage phenomenon but continues to depress satisfaction with advancement after entry, which is why Section 10 recommends bias auditing as an ongoing process rather than a one-time recruitment fix.

9.4 Objective 4 / Hypothesis 4 — Policy and Self-Efficacy

Organisational Feminist Policy scored lowest of all constructs ($M = 2.6$), while Self-Efficacy sat moderately above the midpoint ($M = 3.4$); the positive correlation ($r = .44$) supports H4. The gap between low perceived policy support and moderate self-efficacy suggests individual resilience is currently compensating for weak organisational infrastructure — a pattern that, if found in real data, would be a caution against attributing women's AI-career persistence to organisations doing enough, when the illustrative data suggest the opposite.

Taken together, the four illustrative results form a coherent, internally consistent story precisely because they were constructed to test the framework in Section 2 rather than sampled from a real population. A genuine study using this design might replicate this pattern, contradict it, or (most likely) find a messier

mixture — which is itself a reason the demonstration design in Section 7 should be treated as a starting point for primary research rather than a conclusion.

10. Policy and Practice Recommendations

Recommendation 1 — Publish role-level, not just workforce-level, gender data. Aggregate “% women in AI” figures obscure the leadership drop-off documented in Figure 1; organisations should report the same breakdown internally.

Recommendation 2 — Fund structured mentorship with accountability, not informal pairing alone, since the illustrative H2 result suggests mentorship access lags persistence intention rather than driving it.

Recommendation 3 — Treat algorithmic bias auditing as a retention issue, embedding it in post-hire review cycles rather than only in hiring pipelines.

Recommendation 4 — Adopt feminist-informed policy benchmarks (pay-equity audits, transparent promotion criteria, flexible work) and measure them with the same rigor as technical KPIs.

Recommendation 5 — Invest upstream in girls' AI and robotics education, given the UNESCO finding that only around three per cent of female higher-education students currently choose ICT fields.

11. Limitations

This article combines two different evidentiary registers, and the distinction matters. Section 3's statistics are drawn from real, cited, publicly available sources and are subject only to the normal limitations of secondary research — differing methodologies, national coverage, and definitions of “AI role” across the underlying studies. Sections 7 through 9, by contrast, present an illustrative demonstration design; the sample, means, correlation matrix, subgroup breakdown, and qualitative themes are constructed for exposition and must not be cited as empirical findings. A genuine follow-up study would need institutional review, a probability or well-justified purposive sample substantially larger than fifty respondents — and considerably larger still to power the intersectional subgroup comparisons proposed in Section 6.4 — validated psychometric instruments, a documented reflexive thematic-analysis procedure for the qualitative component, and pre-registration of the hypotheses tested here. The feminist methodological commitments in Section 6 describe how such a study should be conducted; they do not, on their own, substitute for conducting it.

12. Conclusion

The gender gap in AI careers narrows women's access to the fastest-growing, best-paid segment of the global labour market, and it narrows further at every step toward technical and executive leadership. A feminist, technofeminist, and intersectional framework reframes this gap as a structural and technopolitical problem rather than a supply-side pipeline problem alone, implicating workplace culture, algorithmic design, and organisational policy alongside education. The demonstration research design offered here — with its illustrative, clearly labelled results — is intended as a template that researchers, HR teams, or advocacy organisations can adapt and re-run with real respondents, so that claims about mentorship, bias, and policy in AI careers rest on data rather than assertion.

References

1. Collins, P. H. (1990). *Black feminist thought: Knowledge, consciousness, and the politics of empowerment*. Routledge.

2. Faulkner, W. (2001). The technology question in feminism: A view from feminist technology studies. *Women's Studies International Forum*, 24(1), 79–95.
3. Girl Scouts of the USA, cited in World Economic Forum. (2021). 5 ways to increase the number of women working in AI. <https://www.weforum.org/stories/2021/08/5-ways-increase-women-working-ai/>
4. Haraway, D. (1988). Situated knowledges: The science question in feminism and the privilege of partial perspective. *Feminist Studies*, 14(3), 575–599.
5. Harding, S. (1991). *Whose science? Whose knowledge? Thinking from women's lives*. Cornell University Press.
6. Idemudia, I. (n.d.). Gender, workforce and artificial intelligence. Robert Gordon University Open Access Institutional Repository. <https://rgu-repository.worktribe.com/OutputFile/2098632>
7. Reinharz, S. (1992). *Feminist methods in social research*. Oxford University Press.
8. LinkedIn Economic Graph. (2026). The triple penalty: Women in AI leadership, AI jobs, and AI companies [Reported in]. *AfroTech*. <https://afrotech.com/13-of-ai-executives-are-women-heres-why-thats-a-problem>
9. LinkedIn Economic Graph. (2026). Women made up just 26% of AI hires in 2025 [Reported in]. *Allwork.Space*. <https://allwork.space/2026/08/women-made-up-just-26-of-ai-hires-in-2025-linkedin-finds/>
10. Take The Lead. (2025). The good, the bad, and the promising for women in generative AI workforce. <https://taketheleadwomen.com/blog/the-good-the-bad-and-the-promising-for-women-in-generative-ai-workforce>
11. University of Sussex. (2024). Feminism confronts AI: The gender relations of digitalisation. <https://www.sussex.ac.uk/broadcast/read/63624>
12. Voutyrakou, D. A., & Skordoulis, C. (2025). Using AI tools to enhance educational robotics to bridge the gender gap in STEM. *Education Sciences*, 15(6), 711.
13. Wajcman, J. (2004). *TechnoFeminism*. Polity Press.
14. WomenTech Network. (n.d.). How can we leverage tech to break gender barriers in STEM fields? <https://www.womentech.net/how-to/how-can-we-leverage-tech-break-gender-barriers-in-stem-fields>
15. World Economic Forum. (2025). Agenda dialogues: Women in AI [Podcast]. <https://www.weforum.org/podcasts/agenda-dialogues/episodes/women-in-ai>
16. EducationHQ. (2026). Digital disadvantage: AI threatens to widen gender gap for regional girls. <https://educationhq.com/news/digital-disadvantage-ai-threatens-to-widen-gender-gap-for-regional-girls-195512/>