

October 2026

ARE WOMEN BEING LEFT OUT OF AI?

New evidence on who builds, uses, trusts,
shapes, and gets heard in the AI era

Principal Investigators:

Nancy Gibbs

Lombard Director of the Shorenstein Center, Edward
R. Murrow Professor of the Practice of Press, Politics
and Public Policy at Harvard Kennedy School

Matthew Baum

Marvin Kalb Professor of
Global Communications at
Harvard Kennedy School

Report Authors:

Laura Manley

Executive Director

Brian Friedberg

Senior Researcher

Allison Wan

Senior Data Scientist



HARVARD Kennedy School

SHORENSTEIN CENTER

on Media, Politics and Public Policy

TABLE OF CONTENTS

Executive Summary	3
The Pattern Across the AI Ecosystem	5
About This Research	6
Key Findings	7
Usage	7
Knowledge and Awareness	9
Media	10
Attitudes, trust, and barriers to use	12
Workplace Training, Mentorship and Encouragement	16
Representation, Bias and Discrimination	17
What the Findings Suggest	20
Acknowledgements	21
Appendix A: Supplementary Charts and Tables	22
Appendix B: Methodology	41
Appendix D: Literature Review	44
Works Cited	45

EXECUTIVE SUMMARY

Generative artificial intelligence (AI), technology that can create new content such as text, images, audio, and code, has moved from science fiction to the center of daily life in a matter of years. The public is still forming its sense of who builds it, who uses it, who understands it, whose needs it reflects, and who are the experts to trust. Across each of these dimensions, gender differences emerge.

This research establishes an empirical baseline for how people are experiencing generative AI in the United States. It combines Civic Health of Institutions Project (CHIP50) survey data from more than 5,000 U.S. adults, measures of online browsing behavior from the National Internet Observatory (NIO), a meta-analysis of more than 78,000 U.S. news articles about AI, and existing academic and policy research.

Taken together, the findings show a consistent pattern: women use AI less often, report less support for using it at work, appear less often in public discussion of AI, and are perceived as less well represented in its development and outputs. Key findings include:

- ▶ In our survey, women are more likely than men to have never used AI and less likely to use it regularly. Desktop browsing data show no statistically significant difference between women and men in having ever used AI. In the same data, men who use AI spend a larger share of their browsing time on AI tools than women who use AI (4.6% versus 2.9%).
- ▶ Among employed adults, women are less likely than men to report access to AI training, mentorship, and workplace encouragement to use AI. These forms of support are associated with greater AI use, confidence in using AI, and perceived value among both women and men.
- ▶ Across genders, survey respondents report low trust in AI, with women reporting lower trust than men. Women are more likely to associate AI with concern, worry, and anxiety, while men are more likely to associate it with curiosity, optimism, and excitement.
- ▶ Women are also less visible in news about AI. In our sample, men were named in AI headlines 14.6 times as often as women and cited as experts 4.3 times as often.
- ▶ Across genders, respondents say men's needs and perspectives are better represented in AI development and AI-generated content, while few often see women as expert voices in AI news.

These gaps do not mean that higher AI use or greater trust should be goals in themselves. The question is whether women have equal access to understand and use these technologies in a way that adds value to their lives, equal influence over how they are developed and deployed, and equal visibility and authority in the public conversation surrounding them.

As AI becomes more embedded in work, information, education, healthcare, and everyday decision-making, existing disparities may become more pervasive and consequential. These findings provide a baseline for identifying where those disparities are emerging and where action is needed.

ARE WOMEN BEING LEFT OUT OF AI?

THE PATTERN ACROSS THE AI ECOSYSTEM

The gender gap in AI is often framed as a workforce issue, with women underrepresented in technical and leadership roles. Our research shows that gender disparities also extend beyond the industry itself.

Development

Women remain **underrepresented in AI leadership and technical roles**, which can limit their participation in decisions about how these technologies are built.

Access and support

Women are **less likely than men to report workplace encouragement** to use AI or access to AI training and mentorship.

Use

In our survey, women are more likely than men to have never used AI and **less likely to use it regularly**. In desktop browsing data, men who use AI spend a larger share of their browsing time on AI tools than women who use AI.

Trust and understanding

Respondents across genders report low trust in AI, with **women reporting lower trust**. Women are also more likely to say they are unsure about AI and to associate it with concern and anxiety.

Representation

Respondents across genders perceive men's **needs and perspectives as better represented** in AI development and AI generated content.

Expertise and visibility

In news about AI, women are **far less likely than men to be featured** as leaders or expert sources. Men also dominate AI headlines and the experts cited within coverage.

These gaps extend beyond the AI workforce, appearing in workplace support, AI use, trust and uncertainty, representation, and visibility in news coverage.

ABOUT THIS RESEARCH

Artificial intelligence is changing how people work, seek information, and make decisions daily. Evidence points to gender differences in how people engage with AI and in how AI represents and treats women.

The research compares how women and men use, understand, and trust AI, and what holds them back from using it more. It also covers workplace conditions linked to use, representation in AI, perceived and documented bias, and how AI news portrays women and men. Most survey respondents identified as men or women, so our analysis centers on these two groups, though we recognize that not everyone identifies within a gender binary.

The analysis draws on several sources of evidence,

- ▶ CHIP50 survey data from 5,403 U.S. adults collected in July 2026.
- ▶ Anonymized NIO browsing data from 5,513 U.S. adults collected from July 2024 to December 2025 and 1,951 U.S. adults from July 13 to August 10, 2026
- ▶ An analysis of 78,542 U.S. news articles about AI published between July 2025 and June 2026.
- ▶ Academic and policy literature providing context for the findings.

These findings are limited to respondents who self-identify as men or women. Small sample sizes prevent us from drawing reliable insights about individuals who identify outside this gender binary, highlighting an important area for future research. Unless specified, reported gender differences between men and women are statistically significant.

Overall, this report highlights where gender differences are largest and provides recommendations for action.

KEY FINDINGS


USAGE

Women report using AI less than men

Usage measures whether and how often people use AI tools across personal and work settings. In our July 2026 survey, women were more likely than men to say they have never used AI (29% versus 22%). Men were more likely than women to say they use AI occasionally (27% versus 24%) or regularly (25% versus 16%), or that they create AI workflows or automations (5% versus 3%). Women were more likely to say they have tried AI but rarely use it (28% versus 21%). The gap holds across tools, tasks, and personal and work contexts (Figures A1 – A3). Usage is higher among younger, more educated, and higher-income adults, as well as students, employed adults, urban residents, and Asian Americans (Figure A4). The gender gap is most pronounced among those with higher education and income, and is minimal among lower socioeconomic groups, where AI use is lower.

Women report lower levels of AI usage than men

% of respondents who describe their experience with AI tools as:



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.


Unweighted sample size: 2,606 women, 2,431 men. 95% margin of error: ± 0.7 – 1.9 percentage points for women, ± 0.9 – 1.9 for men.

All differences between women and men are statistically significant.

In desktop browsing data from July 2024 to December 2025, we found no statistically significant difference between women and men in having ever used AI.

Similar shares of women and men have used AI at least once

% of participants who had used AI at least once by each month



Source: National Internet Observatory browsing data, July 2024–December 2025. Weighted to U.S. adults with internet access on a desktop or non-mobile device.


Unweighted sample size: 5,513 participants. 95% margin of error: ± 2.6 – 3.7 percentage points for women, ± 2.2 – 4.0 for men.

Hollow points: the difference between women and men is not statistically significant.

In desktop browsing data from July 2024 to December 2025, men who used AI in a given month spent a larger share of their browsing time on AI tools than women who used AI (4.6% versus 2.9% across the full period). This difference is statistically significant across the full period but not in every month.

Women spend a lower share of browsing time using AI

% of browsing time spent using AI tools averaged across participants using AI in each month



Source: National Internet Observatory browsing data, July 2024–December 2025. Weighted to U.S. adults with internet access on a desktop or non-mobile device.

Unweighted sample size: 5,513 participants. 95% margin of error: ± 0.9 – 2.6 percentage points for women, ± 1.1 – 1.9 for men.

Hollow points: the difference between women and men is not statistically significant.

KNOWLEDGE AND AWARENESS

Women express more uncertainty about AI, but that does not necessarily mean they know less

Knowledge and awareness capture how much people know about AI, how closely they follow developments, and how familiar they are with prominent AI tools and figures. Women report paying less attention to AI developments (Figure A5), are less familiar with common AI tools (Figure A6), and express greater uncertainty about AI. Across 12 factual questions (Figure A7), women answered an average of 6.3 correctly, compared with 7.1 among men, with men more likely to answer each question correctly. The gap in the number of incorrect answers is closer (1.6 among women and 2.0 among men), while women selected “not sure” more often (4.0 questions on average, compared with 2.9 among men).

Women also say “not sure” more often than men on opinion-driven AI-related questions (% stating “not sure” to each question listed below, women first, men second):

- ▶ Whether men’s or women’s needs are better represented in AI development: 43% versus 31%.
- ▶ Which AI topics are most associated with women or with men: 38% versus 29%.
- ▶ Whether AI news coverage features mostly men or women: 25% versus 18%.
- ▶ How often they notice women as expert voices: 33% versus 24%.
- ▶ Whether certain groups are treated unfairly or harmed by AI: 42% versus 31%

These patterns suggest greater uncertainty among women, but some of the apparent gender gap in knowledge may also reflect greater confidence, or potentially over-confidence, among men rather than greater knowledge alone. When asked whether they recognized 12 influential people in AI alongside two fabricated names, men were more likely to report recognizing every name including the fabricated ones (Figure A8).

These patterns suggest greater uncertainty among women, but some of the apparent gender gap in knowledge may also reflect greater confidence, or potentially over-confidence, among men rather than greater knowledge alone.

In a similar pattern to usage, gender differences in how much respondents pay attention to developments in AI are largest among people with a graduate degree, between the ages of 25 and 44, living in urban areas, and with household incomes over \$100,000 (Figure A9), which correspond to the demographics likely to be working with, or even developing AI in their professional lives (Sajadieh et al. 2026; Pew Research Center 2026). By contrast, factual knowledge of AI has narrower gender gaps among those with high levels of education and income (Figure A10).

MEDIA

Men overwhelmingly dominate the public conversation about AI

Analysis of AI News Articles

Women are a minority in the AI field itself. World Economic Forum data show that in 2025 about 29% of workers listing AI engineering skills were women, still under a third of the AI talent pool, and that women hold only about 12% of the most senior technology leadership roles (WEF 2025). The people who lead AI companies are overwhelmingly men, and leaders are the people in the field who appear most often in the news.

Across one year of U.S. news coverage about AI (78,542 articles collected through Media Cloud from July 2025 to June 2026), men are named and quoted far more often than women. We analyzed these articles in two ways: headlines and article text. Across the headlines, men are named 14.6 times as often as women.. Among the individuals named in five or more headlines, 219 are men and 34 are women, and men account for 93.6% of all headlines mentions against 6.4% for women. The people named most often are the men who lead AI companies.

Across the headlines, men are named 14.6 times as often as women..

Article text shows the same imbalance. We coded every source an article quoted or paraphrased, 252,233 in total. Just under half, 49.2%, are experts, meaning named individuals cited for their authority, position, or expertise. Among these 124,004 experts, men are cited 4.3 times as often as women. 70.1% are men, 16.2% are women, and 13.6% are people whose gender the article did not identify. Another 42.6% of all sources are organizations quoted as speaking for themselves. Counting those organizations alongside individuals, the 214,515 citations of authority experts break down to 50.1% organizations, 40.5% individual men, and just 9.4% individual women. Women make up fewer than one in ten sources cited as authorities on AI.

Among these 124,004 experts, men are cited 4.3 times as often as women.

Because the men who run AI companies dominate this sample, we tested whether they alone explain the gap. Removing every source cited for a position rather than expertise, including executives, founders, officials, and investors, leaves 16,355 people quoted purely as subject experts. Among them, women are 24.8% of those whose gender was identified, up from 18.8% across all experts whose gender was identified. The gap narrows, but women remain about a quarter of the experts cited for their knowledge, below their roughly 29% share of AI talent.

Demand and Exposure to Information about AI


NIO browsing data shows that AI-related content (i.e. content about AI, as identified using relevant keywords. See Appendix B for more details.) is more often encountered through social media than through news websites. While 41% of panelists visited a news website containing at least one AI-related keyword, 68% encountered an AI-related keyword on social media. These differences partly reflect broader use of these website types: 62% of panelists visited a news website, compared with 76% who visited social media at least once in the data collection period.

The data also show gender differences both in how often men and women visit social media and news websites and how frequently they encounter AI-related content when they do. Similar percentages of men and women saw AI-related content on a news website. However, AI-related keywords appeared in 18% of men's visits to news websites, compared with 12% of women's. For social media, a greater share of women encountered AI-related content compared to men, reflecting women's higher social media use overall. AI-related content appeared at similar rates for men (15%) than women (13%).

Finally, 1.3% of search queries by men involve an AI-related keyword compared to 0.6% of search queries by women. 11% of women have a search query with an AI-related keyword compared to 13% of men. These differences in search are not statistically significant.

Reach of AI-related content across websites types and search queries

% of participants who visited each website type or conducted a search query at least once. Lighter bars show any activity; darker bars show activity involving AI-related content.



Source: National Internet Observatory browsing data, July 13 – August 1, 2026. Weighted to U.S. adults with internet access on a desktop or non-mobile device.

95% margin of error for AI-related content (darker bars): ± 2.6 – 4.1 percentage points for women, ± 2.4 – 3.5 for men.

* The difference between women and men in AI-related content is not statistically significant.

ATTITUDES, TRUST, AND BARRIERS TO USE


Women are more skeptical of AI, while trust is low across the board

Concern and worry rate highest as emotional associations with AI for both genders, and women identify with these feelings more strongly, reporting more fear and anxiety than men. Men rate every positive feeling—curiosity, enthusiasm, optimism, interest, and excitement—higher than women do.

This category covers how much people trust AI, the confidence and value they get from using it, their emotional reactions to it, and what keeps usage low. Concern and worry rate highest as emotional associations with AI for both genders, and women identify with these feelings more strongly, reporting more fear and anxiety than men. Men rate every positive feeling—curiosity, enthusiasm, optimism, interest, and excitement—higher than women do.

Concern and worry are most associated with AI, with women expressing more negative feelings overall

Average rating of how strongly each word describes respondents' feelings about AI, from 1 (not at all) to 10 (very much)



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 2,592–2,605 women, 2,421–2,432 men. 95% margin of error: ± 0.11 – 0.13 points for women, ± 0.12 – 0.13 for men.


Hollow points: the difference between women and men is not statistically significant.

Confidence in use and perceived value of AI follow the same pattern at more moderate levels. Women report lower value (2.0 versus 2.3 on a 4-point scale) and lower confidence (2.5 versus 2.7 on a 4-point scale) across tasks, with “learning something new” and “creative projects” scoring the highest for both genders (see Figures A11 – A12 for specific tasks). Gender gaps are the widest among the highly educated and homemakers (Figures A13 – 14).

Trust in AI is low across the board. No single source is trusted by a majority of respondents to give accurate AI information, and women report lower trust than men on every measure: nearly a quarter of women (23%) trust none of the groups asked about, compared with 15% of men. Technical experts and independent researchers are the most trusted sources of information. Employers and coworkers are the most trusted to use AI responsibly (Figure A15 – A17 for additional trust related batteries).

No group of people is trusted by a majority of respondents for information about AI, with trust even lower among women

% respondents saying they would trust ___ to give accurate and reliable information about AI



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 2,605 women, 2,434 men. 95% margin of error: ± 0.8 – 1.9 percentage points for women, ± 1.2 – 2.1 for men.


All differences between women and men are statistically significant.

Privacy concerns and distrust of AI's outputs are reported as the top two barriers to use for both genders.

This lack of trust and low perceived value is consistent with explanations from both men and women on why they do not use AI more frequently. Privacy concerns and distrust of AI's outputs are reported as the top two barriers to use for both genders, followed by a preference to complete tasks without AI (more common among women) and low perceived value (more common among men). By contrast, expression of fear or judgement, a barrier highlighted in recent literature (Cranney, Delecourt, and Koning 2026), is less common. It ranks near the bottom of stated barriers, chosen by just 5% of women and 6% of men (see Figure A18 for more details). While women report these top barriers at higher rates, men and women rank them similarly in importance. This is in line with broader research that finds concerns of privacy and trustworthiness account for roughly a quarter of the gap in usage between men and women. Other research finds a trust gap between women and men, particularly around whether companies will protect personal data (Aldasoro et al. 2024; Arbanas et al. 2024; Cranney, Delecourt, and Koning 2026).

Privacy and trust concerns are top reasons for low AI usage

% of respondents who use AI occasionally, rarely, or never that selected the following as barriers to greater usage:



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 2,081 women, 1,663 men. 95% margin of error: ± 0.8 –2.2 percentage points for women, ± 0.9 –2.4 for men.

* The difference between women and men is not statistically significant.


WORKPLACE TRAINING, MENTORSHIP AND ENCOURAGEMENT

Women have less access to AI resources and support at work

The workplace dimension covers how professional conditions are associated with AI use and value. Overall, we find that being in a work environment that encourages AI use, access to training, and mentorship is associated with greater usage, confidence in use, and perceived value across genders. However, among employed respondents, women are less likely to be in workplaces that offer these supports (Figures A19 – A21). Additionally, gender gaps often persist even among individuals with greater training, mentorship, and encouragement.

Workplace training and support associated with greater usage, confidence, and perceived value, but gender gaps persist

Average usage, confidence, and perceived value by access to workplace training, mentorship, and encouragement among employed respondents



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 147–729 women, 143–747 men. 95% margin of error (points): ± 0.05 – 0.16 for confidence, ± 0.07 – 0.17 for overall AI usage, ± 0.06 – 0.15 for perceived value, ± 0.12 – 0.25 for work usage.

Hollow points: the difference between women and men is not statistically significant.

Women are also less likely to be employed in the first place (ILO 2026). Among those who are, women are more concentrated in jobs whose tasks are most exposed to automation by generative AI. Globally, 4.7% of women's employment falls in the highest-exposure occupations, against 2.4% of men's, and in high-income countries the gap is wider still, 9.6% of women's jobs versus 3.5% of men's (Gmyrek et al 2025). This reflects occupational segregation, with women overrepresented in clerical and administrative roles that generative AI can most readily perform. Separately, a study of researchers found men's productivity gains after ChatGPT's release outpaced women's by more than six percentage points, alongside workplace data showing far more men than women had even tried the tool, 65% versus 47% (Tang et al, 2025).

REPRESENTATION, BIAS AND DISCRIMINATION


Women are less represented in both AI systems and the expertise surrounding them

A consensus emerged across genders once “not sure” answers were set aside: men’s perspectives and needs are better represented in AI development and in the content AI produces, men are more associated with AI news and online content, and women’s visibility as expert voices is lower.

This category covers perceptions of women's representation in AI, whether AI treats any groups unfairly, and documented bias in AI systems themselves. A consensus emerged across genders once “not sure” answers were set aside: men's perspectives and needs are better represented in AI development and in the content AI produces, men are more associated with AI news and online content, and women's visibility as expert voices is lower. Women are more associated with arts, healthcare, education, and social impact. Men are more associated with business leadership, engineering, government and policy, and national security (Figure A22).

Across genders, respondents see men as more represented and visible in AI

% of respondents perceiving greater representation of women, men, or neither, and reporting how often they notice women as expert voices in AI news



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,411–1,759 women, 1,636–1,859 men. 95% margin of error: ± 1.2 –2.8 percentage points for women, ± 1.3 –2.6 for men.

"Not sure" responses are excluded.

On name recognition, even the most recognized woman in AI, Daniela Amodei (49%), was slightly less recognized than the least recognized men, Jensen Huang and Dario Amodei (50% each). Fei-Fei Li and Mira Murati were each recognized by 46% of respondents, the same share that reported recognizing two fabricated, common-sounding names, Susan Millner and Craig Dalton. Recognition was highest for male technology leaders whose public profiles extend well beyond AI, including Tim Cook (66%), Sundar Pichai (53%), and Satya Nadella (52%), and higher still for Elon Musk, Mark Zuckerberg, and Jeff Bezos (all above 85%). Sam Altman, the most recognized leader tied specifically to AI, was named 64% of the time. See Figure A6 for a breakdown by respondent gender. This is consistent with NIO browsing data showing that 17%, 11%, 7%, and 0.3% of individuals visited a webpage with "Sam Altman," "Jensen Huang," "Dario Amodei," or "Satya Nadella" appearing on it respectively. By contrast, 1.6%, 1.3%, and 0.1% of respondents saw "Daniella Amodei," "Mira Murati," or "Fei-Fei Li" respectively.

On unfair treatment specifically, women are more likely to be unsure whether AI treats any group unfairly (42% versus 31% of men), while men are more likely to say that no group is.

Women and men report the presence of AI bias at about equal rates, and both name Black people and low-income people as the groups treated unfairly most often. On gender specifically, each gender is somewhat more likely to say its own gender is the one treated unfairly (Figure A23). While similar shares of women (58%) and men (59%) perceive AI bias toward some groups, and 47% of women and 43% of men belong to a demographic group they believe AI treats unfairly, only 11% of women and 10% of men report personally experiencing AI-related harm (Figure A24). This suggests that perceptions of AI bias may be shaped more by social networks and media coverage than by personal experience.

WHAT THE FINDINGS SUGGEST

Workplace access

Women are less likely to report access to AI training, mentorship, and encouragement at work even as these forms of support are associated with greater confidence, perceived value, and use among women and men. Employers, industry groups, and professional associations should make sure women have equal access to the information and support they need to decide when and how to use AI in their work.

Low trust and negative attitudes

Across both genders, trust in who develops AI, and who provides information about AI is low. Privacy and trust concerns are the most common barriers to using AI more, and the feelings most strongly associated with AI are concern and worry. Research should continue investigating the reasons behind these concerns, while journalists, developers and policymakers should work to build trust with greater transparency, stronger privacy protections, and the provision of credible information.

Women in AI development

Women remain underrepresented in technical and leadership roles that shape AI systems. The people in such roles make product and safety decisions and help decide whose needs are considered during development. At best, underrepresentation may mean missed opportunities for value creation among certain groups. At worst, it may cause direct harm, reinforce existing social inequities, sow distrust, and perpetuate pipeline issues that limit who enters and advances in the field. Existing research also documents gender bias in images and text created by AI, hiring systems, healthcare applications, and other uses.

Who gets heard as an expert

Men dominate AI news headlines and expert citations. While this reflects, in part, the prominence of male founders and executives. In expert citations, the imbalance continues after those leaders are excluded from the analysis. Journalists, researchers, companies, and conveners can widen the pool of experts they feature.

Better measurement

AI is changing quickly. Continued research should track these gaps over time and examine differences among women across demographic groups that national averages can conceal. It should also examine how people encounter information about AI beyond traditional news, how people interact directly with AI, and how these experiences shape attitudes and behavior.

Acknowledgements

The research team included co-Principal Investigators Nancy Gibbs, Edward R. Murrow Professor of the Practice of Press, Politics and Public Policy at Harvard Kennedy School, and Matthew Baum, Marvin Kalb Professor of Global Communications, Harvard Kennedy School. Lead researchers Laura Manley, Executive Director; Allison Wan, Senior Data Scientist; and Brian Friedberg, Senior Researcher co-authored the report. Special thank you to David Lazer, University Distinguished Professor of Political Science and Computer Sciences at Northeastern University, for his input and guidance. Thanks to Kathy Pham and Kasia Chmielinski for their advice and input on this report.

This research and report were funded by Pivotal. Founded by Melinda French Gates, Pivotal is a group of organizations working to accelerate the pace of social progress for women and young people in the U.S. and around the world.

APPENDIX A


SUPPLEMENTARY CHARTS AND TABLES

The charts below supplement the findings in the main report with additional details on usage, confidence, knowledge, attitudes, trust, workplace support, media associations, and perceived bias.

Figure A1

Levels of work-related AI usage

% of employed respondents answering "How often do you use AI for your work?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,152 women, 1,427 men. 95% margin of error: ± 1.4 – 3.1 percentage points for women, ± 1.3 – 2.7 for men.

* The difference between women and men is not statistically significant.

Figure A2**AI usage by task**

Average usage score (1 = Never, 4 = Often) among respondents using AI answering "How often have you used AI tools for each of the following tasks, over the past 3 months?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,843–1,852 women, 1,909–1,920 men. 95% margin of error: ± 0.05 points for women, ± 0.05 for men.

All differences between women and men are statistically significant.

Figure A3**AI usage by tool**

Average usage score (1 = Never, 7 = Several times a day) among respondents answering "How often do you use the following AI tools?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 24–2,241 women, 39–2,152 men. 95% margin of error: ± 0.08 – 0.79 points for women, ± 0.09 – 0.53 for men.

Hollow points: the difference between women and men is not statistically significant.

Figure A4**AI usage by demographics**

Average usage score (1 = "I have never used AI tools"; 5 = "I create or customize AI workflows or automations") among respondents answering "Which best describes your experience with AI tools?"



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 10–1,810 women, 10–1,727 men. 95% margin of error: ± 0.06 –0.84 points for women, ± 0.06 –0.60 for men.


Race and ethnicity are check-all-that-apply, so groups overlap and do not add to 100%.

Hollow points: the difference between women and men is not statistically significant.

Figure A5

Attention paid to developments in AI

% of respondents answering "How much attention do you pay to developments in the area of artificial intelligence (AI)?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.
Unweighted sample size: 2,608 women, 2,433 men. 95% margin of error: ± 1.1 –2.0 percentage points for women, ± 1.6 –2.0 for men.
All differences between women and men are statistically significant.

Figure A6

Awareness of AI tools

% of respondents answering "Yes" to "Have you heard of the following AI tools?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.
Unweighted sample size: 2,582–2,606 women, 2,399–2,429 men. 95% margin of error: ± 0.9 –2.0 percentage points for women, ± 1.2 –2.1 for men.
Hollow points: the difference between women and men is not statistically significant.

Figure A7

Accuracy and uncertainty of responses to questions assessing knowledge about AI

% of respondents saying each statement is accurate, inaccurate, or that they are not sure



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 2,595–2,601 women, 2,422–2,430 men. 95% margin of error: ± 1.1 – 2.0 percentage points for women, ± 1.2 – 2.1 for men.

† Statement is inaccurate.


Hollow points: the difference between women and men is not statistically significant.

Figure A8

Name recognition of individuals developing AI

% of respondents saying they know who ___ is

Note: Susan Millner and Craig Dalton are fabricated names



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.


Unweighted sample size: 2,594–2,608 women, 2,422–2,432 men. 95% margin of error: ± 0.8 –2.0 percentage points for women, ± 0.5 –2.1 for men.

All differences between women and men are statistically significant.

Figure A9

Attention to AI by demographics

Average attention score (1 = None, 4 = A lot) among respondents answering "How much attention do you pay to developments in the area of artificial intelligence (AI)?"



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 10–1,812 women, 10–1,728 men. 95% margin of error: ± 0.04 – 0.54 points for women, ± 0.05 – 0.59 for men.


Race and ethnicity are check-all-that-apply, so groups overlap and do not add to 100%.

Hollow points: the difference between women and men is not statistically significant.

Figure A10

AI knowledge by demographics

Average number of 12 factual statements about AI answered correctly



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.


Unweighted sample size: 10–1,812 women, 10–1,728 men. 95% margin of error: ± 0.19 – 2.47 points for women, ± 0.18 – 2.05 for men.

Race and ethnicity are check-all-that-apply, so groups overlap and do not add to 100%.

Hollow points: the difference between women and men is not statistically significant.

Figure A11**Perceived value of using AI by task**

Average value score (1 = Not valuable at all, 4 = Very valuable) among respondents using AI answering "How valuable have you personally found AI tools for each of the following tasks?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,834–1,850 women, 1,902–1,915 men. 95% margin of error: ± 0.05 – 0.06 points for women, ± 0.05 – 0.06 for men.

All differences between women and men are statistically significant.

Figure A12**Confidence using AI by task**

Average confidence score (1 = Not at all confident, 4 = Very confident) among respondents using AI answering "How confident are you in your ability to use AI tools effectively for each of the below tasks?"



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.


Unweighted sample size: 1,834–1,848 women, 1,903–1,916 men. 95% margin of error: ± 0.05 points for women, ± 0.05 for men.

All differences between women and men are statistically significant.

Figure A13

Perceived value of AI by demographics

Average value score across tasks (1 = Not valuable at all, 4 = Very valuable) among respondents using AI



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.


Unweighted sample size: 8–1,250 women, 9–1,346 men. 95% margin of error: ± 0.05 – 0.40 points for women, ± 0.05 – 0.63 for men.

Race and ethnicity are check-all-that-apply, so groups overlap and do not add to 100%.

Hollow points: the difference between women and men is not statistically significant.

Figure A14**Confidence using AI by demographics**

Average confidence score across tasks (1 = Not at all confident, 4 = Very confident) among respondents using AI



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 8–1,245 women, 9–1,346 men. 95% margin of error: ± 0.05 – 0.59 points for women, ± 0.05 – 0.39 for men.


Race and ethnicity are check-all-that-apply, so groups overlap and do not add to 100%.

Hollow points: the difference between women and men is not statistically significant.

Figure A15**Trust in individuals developing AI**

Average trust score (1 = Not at all, 4 = A great deal) among respondents who recognized each person answering "How much would you trust the following people to develop AI responsibly?"

Note: Susan Millner and Craig Dalton are fabricated names



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.


Unweighted sample size: 795–2,236 women, 1,064–2,238 men. 95% margin of error: ± 0.05 – 0.08 points for women, ± 0.05 – 0.07 for men.

Includes only respondents who recognized each name; "Not sure" is excluded.

All differences between women and men are statistically significant.

Figure A16**Trust in AI outputs**

Average trust score (1 = Not at all, 4 = A great deal) among respondents answering "How much do you trust AI tools to give you accurate information about..."




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 2,592–2,603 women, 2,427–2,433 men. 95% margin of error: ± 0.04 points for women, ± 0.04 for men.

All differences between women and men are statistically significant.

Figure A17**Trust in organizations to use AI responsibly**

Average trust score (1 = Not at all, 4 = A great deal) among respondents answering "How much do you trust the following organizations to use AI responsibly?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,151–2,603 women, 1,424–2,432 men. 95% margin of error: ± 0.04 – 0.06 points for women, ± 0.04 – 0.05 for men.

All differences between women and men are statistically significant.

Figure A18**Judgement and negative attitudes towards using AI at work**

Average agreement (1 = Strongly disagree, 5 = Strongly agree) among employed respondents to the following statements:




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,142–1,152 women, 1,423–1,427 men. 95% margin of error: ± 0.08 – 0.09 points for women, ± 0.07 – 0.08 for men.

Hollow points: the difference between women and men is not statistically significant.

Figure A19**Workplace encouragement of AI usage**

% of employed respondents answering "How would you characterize your workplace with respect to AI use?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,150 women, 1,426 men. 95% margin of error: ± 1.4 – 3.1 percentage points for women, ± 1.3 – 2.8 for men.

* The difference between women and men is not statistically significant.

Figure A20**Access to AI mentorship**

% of employed respondents answering "I have someone in my life who has shown me how to get useful results from AI tools."




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,151 women, 1,423 men. 95% margin of error: ± 1.9 – 2.8 percentage points for women, ± 1.7 – 2.5 for men.

* The difference between women and men is not statistically significant.

Figure A21**Workplace AI training**

% of employed respondents answering "How would you describe AI training in your workplace?"




Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,151 women, 1,426 men. 95% margin of error: ± 1.8 –3.0 percentage points for women, ± 1.9 –2.8 for men.

All differences between women and men are statistically significant.

Figure A22**Topics associated with women and men in AI**

% of respondents associating each topic with men or women, by respondent gender



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 2,593–2,603 women, 2,418–2,426 men. 95% margin of error: ± 0.9 –1.9 percentage points for women, ± 1.0 –2.0 for men.


Grey bars average the questions about women and about men.

* The difference between women and men is not statistically significant in at least one panel.

Figure A23

Groups perceived as treated unfairly or inaccurately by AI

% of respondents selecting the following answering "Which groups of people, if any, are treated unfairly or inaccurately by AI tools or automated systems?"



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.


Unweighted sample size: 2,604 women, 2,432 men. 95% margin of error: ± 0.5 –2.0 percentage points for women, ± 0.5 –2.0 for men.

Hollow points: the difference between women and men is not statistically significant.

Figure A24

Perceptions and experiences of AI bias by gender

% of respondents perceiving AI bias or reporting personal unfair treatment



Source: CHIP50 survey of U.S. adults, July 10–16, 2026, weighted to the U.S. population.

Unweighted sample size: 1,531–2,604 women, 1,700–2,432 men. 95% margin of error: ± 1.2 –2.7 percentage points for women, ± 1.2 –2.5 for men.

Own demographic group is based on gender, race and ethnicity (check-all-that-apply), age, or income.

Hollow points: the difference between women and men is not statistically significant.

APPENDIX B

METHODOLOGY

CHIP50 methods

From July 10, 2026 to July 16, 2026 we collected a total of 5,403 responses from individuals age 18 and older across all 50 states and the District of Columbia. The surveys were conducted by PureSpectrum through an online, nonprobability sample. We reweighted our data using demographic characteristics to match the U.S. population with respect to 2020 vote choice and turnout, race and ethnicity, age, gender, education, and living in urban, suburban, or rural areas. More information on methodology is available at chip50.org/survey-methodology.

NIO methods

The National Internet Observatory (NIO) is an NSF-funded research infrastructure that collects data from consenting U.S. adults through a browser extension installed on participants' desktop or laptop computers. The extension records browsing activity, which can be linked to participants' self-reported demographic characteristics and is made available to approved researchers through a privacy-preserving data access framework.

The report is based upon two streams of NIO data. The first is desktop browsing activity collected from July 1, 2024 through December 31, 2025 that records when and how long participants spent on each website. We excluded participants with browsing limited to survey websites, less than ten minutes of total browsing time, or fewer than seven distinct browsing days, resulting in a final sample of 5,513 participants.

The second records the websites participants visited through desktop browsing and whether at least one of 772 AI-related keywords (e.g., "AI," "large language model," "chatbot") appeared on each page from July 13 to August 10, 2026. Based on the URL, we are also able to detect search queries and whether an AI-related keyword is included. As these data do not measure time spent on a page, we analyze visits instead. We define a visit as a unique instance of page content viewed within a given day. Returning to the same page without a change in its content (for example, switching tabs and returning) does not count as a new visit. However, when the content of a page changes while the URL remains the same, such as when scrolling through a social media feed, the new content counts as a separate visit. We also excluded participants with browsing limited to survey websites or fewer than two distinct browsing days, resulting in a final sample of 1,951 participants.

Because NIO uses non-probability recruitment, for both data streams we apply NIO-provided weights based on age, gender, income, and education. The weights approximate the population

of U.S. adults with internet access who are not mobile-only internet users, representing approximately 85% of U.S. adults.

AI usage is identified by matching website URLs to a list of known AI tool domains derived from a publicly available dataset of AI tools and platforms curated from AIToolBuzz.com and accessed in October 2025. Classifying websites as news, social media, email are based on lists from Lazer et al. (2026), and refined by the authors.

Several limitations should be considered. NIO participants are recruited through non-probability methods, so weighting may not fully eliminate selection bias. Participants can opt out of tracking through incognito browsing, and the extension is not intended for work computers, meaning the data primarily reflect personal desktop browsing rather than total internet use. The data also do not capture activity that occurs outside the browser, including through desktop and mobile applications such as social media apps and coding agents. Finally, the sample is restricted to adults and therefore does not capture browsing by minors.

Media Analysis

Both studies draw on a single Media Cloud pull of United States online AI news published July 2025 through June 2026: 84,422 article URLs, reduced to 78,542 records by keeping English-language text and removing video, podcast, and gallery pages. The headline study analyzed the headline text of these 78,542 records. We extracted person names with spaCy named-entity recognition (model `en_core_web_sm`), keeping entities tagged PERSON and counting each once per headline, which yielded 9,703 distinct name strings. We retained name strings appearing in five or more headlines and removed a fixed list of company and product terms by text match, leaving 414 name strings. We hand-coded each of the 414 as woman, man, or non-person, which yielded 253 coded as people (34 women and 219 men) and 161 non-people (companies, products, and other entities the term list did not catch), which we excluded. The headline shares in Section 1 are counts of headline mentions among the 253 coded people: 366 mentions for women and 5,354 for men, 5,720 in total.

The expert-citation study analyzed the article bodies of the same 78,542 records. We scraped and cleaned each article body, then coded every source it cited with a large language model (Claude Sonnet 5) prompted with a hand-coded schema, in which a source is any person or organization the article quotes or paraphrases as providing information, a statement, or expertise. Each source received a type (expert, a named individual cited for authority, position, or expertise; press statement, an organization cited as making a statement; experiential, a person cited for lived experience; or none), a gender (woman, man, or unspecified, inferred from pronouns, gendered titles, or clearly gendered given names), and a named-or-generic status. Scraping succeeded for 67,199 articles (85.6 percent); coding covered 63,585 of these

(94.6 percent of those scraped), of which 62,072 carried at least one cited source, producing 252,233 coded sources.

APPENDIX D

LITERATURE REVIEW

A growing body of research documents gender differences in who builds AI, how AI systems represent and treat women, and how women engage with these tools. It provides the backdrop for this report's findings.

Who builds AI. Women hold about a fifth of AI roles worldwide and under 15 percent of senior executive positions in the field (Pal et al. 2024). West, Whittaker, and Crawford (2019) describe this concentration as a diversity crisis, in which demographically narrow teams build systems that carry their assumptions. Labor data and reporting alike point to leadership and research ranks that remain heavily male.

Bias in AI systems. Documented bias runs stronger in images than in text. Image generators amplify occupational gender stereotypes (Guilbeault et al. 2024), depict women as younger than men in higher-status roles (Guilbeault, Delecourt, and Srinivasa Desikan 2025), and render women with more submissive expressions (Sun et al. 2024) and lower facial prominence (Weinmann et al. 2026). Facial-analysis systems misclassify darker-skinned women far more often than lighter-skinned men (Buolamwini and Gebru 2018). In language, models assign stereotypically gendered occupations three to six times more often than labor statistics warrant (Kotek, Dockum, and Sun 2023), cast men as agentic and women as communal in generated reference letters (Wan et al. 2023), and produce sexist content in roughly a fifth of gendered prompts (UNESCO and IRCAI 2024). A systematic review attributes this bias about equally to training data and model design (Hall and Ellis 2023). Hiring results are mixed: some audits find models tilt toward men for higher-status technical roles, though the effect is marginal (Armstrong et al. 2024), while others find models more likely to select women yet recommend lower pay (Gerszberg 2024). Bias extends into medicine, including diagnostic skew by gender and a default to male doctors in generated images (Zack et al. 2024; Currie, Chandra, and Kiat 2024).

Work and economic exposure. Occupations most exposed to generative AI cover nearly twice the share of female as male employment, 4.7 percent versus 2.4 percent, a gap labor researchers tie to occupational segregation (Butt et al. 2026; Gmyrek et al. 2025). In academia, men's measured productivity gains after the release of ChatGPT outpaced women's (Tang et al. 2025).

Adoption, knowledge, and trust. A synthesis of global evidence finds a persistent adoption gap favoring men, holding even within the same occupation or firm (Cranney, Delecourt, and Koning 2026). Differences in self-assessed knowledge explain most of that gap (Aldasoro et al. 2024). Recent survey data show men and women are about equally likely to use chatbots, while women remain more skeptical about AI's impact (Pew Research Center 2026).

Safety and gender-based harm. Sexualized deepfakes fall overwhelmingly on women. An industry scan cited by UN Women (2025) put the share of deepfake-pornography subjects who are women at about 99 percent. Victim gender shapes how the public judges this abuse (Eaton et al. 2026), and scoping reviews argue that understanding and countering it calls for feminist and psychological perspectives that current research still lacks (Lazard et al. 2025).

Works Cited

- Aldasoro, Iñaki, Olivier Armantier, Sebastian Doerr, Leonardo Gambacorta, and Tommaso Oliviero. 2024. "The Gen AI Gender Gap." *Economics Letters* 241: 111814.
- Arbanas, Jana, Paul H. Silverglate, Susanne Hupfer, Jeff Loucks, Prashant Raman, and Michael Steinhart. 2024. "Women and Generative AI: The Adoption Gap is Closing Fast, but a Trust Gap Persists." Technical Report, Deloitte Insights, TMT Predictions 2025.
- Armstrong, Lena, Abbey Liu, Stephen MacNeil, and Danaë Metaxa. 2024. "The Silicon Ceiling: Auditing GPT's Race and Gender Biases in Hiring." In *Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '24)*. New York: ACM.
- Buolamwini, Joy, and Timnit Gebru. 2018. "Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification." *Proceedings of Machine Learning Research* 81: 77–91.
- Butt, Anam Parvez, Emanuela Pozzan, Janine Berg, Pawel Gmyrek, and Souleima El Achkar. 2026. *Gen AI, Occupational Segregation and Gender Equality in the World of Work*. Geneva: International Labour Organization.
- Cranney, Katelyn, Solène Delecourt, and Rembrand Koning. 2026. "Global Evidence on Gender Gaps and Generative AI Over Time." Harvard Business School Working Paper 25-023. Boston: Harvard Business School.
- Currie, Geoffrey, Christina Chandra, and Hosen Kiat. 2024. "Gender Bias in Text-to-Image Generative Artificial Intelligence When Representing Cardiologists." *Information* 15 (10): 594.
- Eaton, Asia A., Adrian J. Scott, Asher Flynn, and Anastasia Powell. 2026. "Perceptions of Sexualized Deepfake Abuse Across Three Nations." *Computers in Human Behavior* 177: 108899.
- Gerszberg, Nina. 2024. "Quantifying Gender Bias in Large Language Models: When ChatGPT Becomes a Hiring Manager." Master's thesis, Massachusetts Institute of Technology.
- Gmyrek, Pawel, Janine Berg, Karol Kamiński, et al. 2025. "Generative AI and Jobs: A Refined Global Index of Occupational Exposure." ILO Working Paper 140. Geneva: International Labour Organization.
- Guilbeault, Douglas, Solène Delecourt, Tasker Hull, Bhargav Srinivasa Desikan, Mark Chu, and Ethan Nadler. 2024. "Online Images Amplify Gender Bias." *Nature* 626 (8001): 1049–55.
- Guilbeault, Douglas, Solène Delecourt, and Bhargav Srinivasa Desikan. 2025. "Age and Gender Distortion in Online Media and Large Language Models." *Nature* 646 (8087): 1129–37.
- Hall, Paula, and Debbie Ellis. 2023. "A Systematic Review of Socio-Technical Gender Bias in AI Algorithms." *Online Information Review* 47 (7): 1264–79.

- Kotek, Hadas, Rikker Dockum, and David Sun. 2023. "Gender Bias and Stereotypes in Large Language Models." In *Proceedings of the ACM Collective Intelligence Conference (CI '23)*, 12–24. New York: ACM.
- Lazard, Lisa, Rose Capdevila, Emma L. Turley, Kim Gilfoyle, and Nikoletta Stavropoulou. 2025. "Deepfake Technology and Gender-Based Violence: A Scoping Review." *Trauma, Violence, & Abuse*.
- Lazer, David, Burak Ozturan, Christo Wilson, Dave Choffnes, Cassidy Waldrip, Hsiu Chi Liu, and John Wihbey. 2026. *How Americans Spend their Time Online*. https://osf.io/preprints/socarxiv/3fhe7_v1
- Pal, Siddhi, Ruggero Marino Lazzaroni, and Paula Mendoza. 2024. "AI's Missing Link: The Gender Gap in the Talent Pool." *Interface*, October 10, 2024.
- Pew Research Center. 2026. *Americans and AI 2026: Chatbots, Smart Devices and Views on Impact*. Washington, DC: Pew Research Center.
- Sajadieh, Sha, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Lapo Santarlaschi, Juan Pava, Nestor Maslej, Russ Altman, Erik Brynjolfsson, Carla Brodley, Jack Clark, Virginia Dignum, Vipin Kumar, James Landay, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Elham Tabassi, Russell Wald, Toby Walsh, and Dan Weld. 2026. *The AI Index 2026 Annual Report*. Stanford, CA: AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.
- Sun, Luhang, Mian Wei, Yibing Sun, Yoo Ji Suh, Liwei Shen, and Sijia Yang. 2024. "Smiling Women Pitching Down: Auditing Representational and Presentational Gender Biases in Image-Generative AI." *Journal of Computer-Mediated Communication* 29 (1).
- Tang, Chuang, Shaobo Li, Suming Hu, Fue Zeng, and Qianzhou Du. 2025. "Gender Disparities in the Impact of Generative Artificial Intelligence: Evidence from Academia." *PNAS Nexus* 4 (2).
- UN Women. 2025. *How AI Is Exacerbating Technology-Facilitated Violence Against Women and Girls*. New York: UN Women.
- UNESCO and IRCAI. 2024. *Challenging Systematic Prejudices: An Investigation into Bias Against Women and Girls in Large Language Models*. Paris: UNESCO.
- Wan, Yixin, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. "Gender Biases in LLM-Generated Reference Letters." In *Findings of the Association for Computational Linguistics: EMNLP 2023*.
- Weinmann, Hannah, Tanja V. Messingschlager, and Markus Appel. 2026. "Gender Bias in Text-to-Image Generative Artificial Intelligence: Neglect and Stereotypical Presentations Across Three Popular Platforms." *New Media & Society*, 1–23.
- West, Sarah Myers, Meredith Whittaker, and Kate Crawford. 2019. *Discriminating Systems: Gender, Race, and Power in AI*. New York: AI Now Institute.
- Zack, Travis, Eric Lehman, Mirac Suzgun, Jorge A. Rodriguez, Leo Anthony Celi, Judy Gichoya, Dan Jurafsky, et al. 2024. "Assessing the Potential of GPT-4 to Perpetuate Racial and Gender Biases in Health Care: A Model Evaluation Study." *Lancet Digital Health* 6 (1): e12–e22.

ARE WOMEN BEING LEFT OUT OF AI?

New evidence on who builds, uses, trusts,
shapes, and gets heard in the AI era



HARVARD Kennedy School

SHORENSTEIN CENTER

on Media, Politics and Public Policy