Bridging Gender Gaps in Women's Healthcare: The Transformative Potential and Challenges of Artificial Intelligence

By Maya Brown

King's College London Policy Research Centre

Centre for Public Health

Word Count: 2443

Summer 2025 Edition

Introduction

Artificial intelligence (AI) is revolutionising healthcare delivery worldwide, offering unprecedented opportunities to address persistent gender disparities in women's health outcomes. As governments and multilateral institutions increasingly recognise AI's role in achieving Sustainable Development Goal 3 (health and wellbeing), this technology emerges as a critical catalyst for public policy innovation. From reproductive health to chronic disease management, AI-driven solutions promise to enhance diagnostic accuracy, improve treatment personalisation, and expand access to care in underserved regions. However, this technological transformation brings complex ethical and practical challenges that demand urgent policy attention, particularly in light of evolving frameworks like the EU AI Act and the WHO's Global Strategy on Digital Health.

The underrepresentation of women in medical research and the male-dominated nature of many healthcare AI systems risk creating new forms of algorithmic discrimination that could undermine recent policy gains in gender equity. Without deliberate intervention anchored in robust governance frameworks, these technologies may perpetuate rather than eliminate existing inequities. This policy brief comprehensively examines AI's role in women's healthcare through a public policy lens, evaluates current regulatory frameworks, and proposes detailed recommendations to ensure these innovations benefit all women equitably while advancing the broader theme of development and innovation in public policy.

My analysis addresses three core questions:

1. How can AI technologies most effectively address gender disparities in healthcare?
2. What are the primary ethical and operational challenges in implementing AI solutions for women's health?
3. What policy interventions are needed to govern AI development and deployment in this sector?

Background: The Current Landscape of AI in Women's Healthcare

Advancements in Diagnostic and Preventive Care

Recently, there has been remarkable progress in AI applications for women's health. Machine learning algorithms now demonstrate superior performance in detecting breast cancer from mammograms, with some systems achieving radiologist-level accuracy while reducing false negatives by up to 9.4% (McKinney et al., 2020). Similarly, AI-powered ultrasound interpretation tools are transforming prenatal care in resource-limited settings, enabling earlier detection of potentially life-threatening conditions like placenta previa.

In maternal health, predictive analytics platforms integrate data from electronic health records, wearable devices, and social determinants of health to identify women at high risk of complications. For example, the QUiPP app, developed by King's College London, uses machine learning to assess preterm birth risk with 86% accuracy, allowing for timely interventions (Sandall et al., 2021).

These advancements coincide with growing policy interest in AI governance, as seen in the EU's proposed AI Act and the WHO's guidance on AI ethics. However, fundamental disparities in health data collection undermine AI's potential, highlighting the need for policy interventions that address technological and systemic barriers.

Persistent Data Gaps and Representation Challenges

Despite these advances, fundamental disparities in health data collection continue to undermine AI's potential. A 2022 analysis of 71 healthcare AI algorithms found that only 32% reported sex-disaggregated performance metrics (Cirillo et al., 2022). This opacity masks potential gender biases in system performance.

The problem extends to training datasets. For instance, many image recognition systems for skin cancer detection were trained primarily on male patients, leading to higher error rates for female cases (Adamson & Smith, 2018). Similarly, conditions predominantly affecting women, including endometriosis, polycystic ovary syndrome (PCOS), and autoimmune diseases, remain underrepresented in AI research and development.

Regulatory and Implementation Challenges

Current policy frameworks struggle to keep pace with technological developments. The EU AI Act, while groundbreaking in establishing risk categories for AI systems, contains limited specific provisions for healthcare applications or gender equity considerations. Internationally, the WHO's guidance on AI ethics provides important principles but lacks enforcement mechanisms, particularly in resource-limited settings where healthcare disparities are most acute.

AI's Transformative Potential in Key Areas

Reproductive Health Innovations

AI is changing fertility care by introducing data-driven precision to traditionally subjective medical decisions. Predictive algorithms analyse hundreds of clinical and laboratory parameters—from hormone levels to genetic markers—to create customised IVF treatment protocols that improve success rates while minimising unnecessary interventions. Advanced computer vision systems now more consistently evaluate embryo viability than human embryologists, using deep learning to assess morphological features imperceptible to the naked eye. These systems enhance pregnancy outcomes and significantly reduce the health risks associated with multiple gestations. Meanwhile, AI-powered chatbots are transforming patient education, offering round-the-clock, judgment-free support for sensitive family planning decisions while collecting data to improve community health outreach strategies.

Maternal Health Monitoring

In underserved regions where maternal mortality remains unacceptably high, AI acts as a force multiplier for overextended healthcare systems. Compact ultrasound devices equipped with automated interpretation algorithms enable community health workers with minimal training to detect fetal malpresentation or placental abnormalities—conditions that would otherwise require specialist consultation. Predictive analytics now process subtle patterns in blood pressure, proteinuria, and other biomarkers to flag preeclampsia risks months before visible symptoms emerge, creating a critical window for preventive care. Natural language processing tools systematically analyse unstructured clinical notes, social determinants of health data, and patient-reported symptoms to identify vulnerable women who might slip through the cracks of conventional care protocols. These technologies don't replace clinicians but create intelligent safety nets where human resources are scarce.

Chronic Disease Management

The gender gap in chronic disease research and treatment is beginning to close through AI's ability to detect sex-specific disease patterns. For perinatal depression—a condition historically underdiagnosed due to stigma and symptom overlap with typical postpartum changes—AI-driven digital therapeutics deliver personalised cognitive behavioural therapy adapted to hormonal fluctuations and sleep deprivation effects. In cardiovascular health, machine learning models trained on female-specific symptom profiles (often distinct from male "textbook" presentations) can now interpret data from wearable sensors to detect early signs of arrhythmia or microvascular dysfunction in postmenopausal women. Perhaps most transformative is AI's role in decoding menstrual cycle data: By analysing long-term tracking information across diverse populations, algorithms can identify subtle irregularities predictive of PCOS, endometriosis, or thyroid disorders years before clinical diagnosis, enabling proactive management.

Systemic Barriers and Risks

Algorithmic Bias and Its Consequences

The healthcare AI revolution carries dangerous blind spots rooted in systemic inequities. Landmark studies reveal how risk prediction algorithms trained primarily on white, male datasets disastrously underestimate care needs for Black women, a flaw stemming from flawed assumptions that healthcare spending equates to healthcare needs. Voice recognition technologies used in clinical documentation frequently mis-transcribe higher-pitched voices, forcing female clinicians to alter their natural speech patterns and potentially distorting patient records. Dermatology AI tools developed with inadequate representation of darker skin tones show alarming accuracy gaps in detecting melanoma in Black patients, exacerbating existing disparities in cancer outcomes. These aren't technical glitches but manifestations of deeper problems: development teams lacking diversity, testing protocols that ignore intersectional identities, and a troubling tendency to treat algorithmic outputs as objective truth rather than probabilistic suggestions shaped by their training data.

Privacy and Autonomy Concerns

The intimate nature of women's health data creates unique vulnerabilities in the AI era. Many fertility apps operate as data brokers, sharing sensitive information about menstrual cycles, sexual activity, and pregnancy intentions with third parties, often buried in opaque terms of service. Predictive analytics in obstetrics risk pathologising normal pregnancies, as algorithms designed to detect complications may pressure women into unnecessary interventions due to defensive medicine practices. Perhaps most alarmingly, inadequate encryption of gynaecological health data could enable discrimination by employers, insurers, or even law enforcement in regions where reproductive rights are contested. These concerns aren't hypothetical—several cases have already emerged of women being targeted with ads for miscarriage services after discreetly searching pregnancy-related terms or facing insurance premium hikes based on fertility app data.

Implementation Challenges

Even the most well-designed AI tools face adoption barriers reflecting healthcare's structural complexities. Clinicians rightfully resist technologies that position algorithms as replacements rather than aids, fearing erosion of hard-won diagnostic skills and therapeutic relationships. Rural deployment stumbles on practical realities: a brilliant AI-powered ultrasound means little in clinics lacking reliable electricity, let alone broadband for cloud-based analysis. Cultural considerations prove equally crucial—in conservative communities, women may avoid digital health tools that require discussing intimate health details with male family members who control device access. These challenges demand more than technical fixes; they require co-design with end-users, investment in human infrastructure alongside technology, and recognition that AI's value lies in culturally competent care.

Policy Recommendations

Strengthening Regulatory Frameworks

The current regulatory landscape for AI in healthcare lacks the specificity to address gender equity concerns, particularly in women's health. Amending the EU AI Act to include explicit requirements for gender equity would represent a critical step forward. These amendments should mandate that healthcare AI systems undergo rigorous bias audits conducted by independent, multidisciplinary bodies, not just the developers themselves, and the results should be publicly available. Furthermore, post-market surveillance protocols must track how these tools perform across different demographic groups in real-world settings. Lab conditions often fail to reveal disparities that emerge in clinical practice.

For high-risk applications in women's health, such as fertility prediction or pregnancy monitoring, specialised guidance should establish concrete standards. This includes protocols for collecting representative datasets that account for racial, socioeconomic, and biological diversity rather than relying on convenience samples that overrepresent specific populations. Evaluation criteria must go beyond aggregate accuracy metrics to assess disparate impact—whether the algorithm performs consistently well for all groups or exhibits dangerous variations. Crucially, these standards should require meaningful involvement from clinicians and patients throughout the development process, ensuring the technology addresses real needs rather than perceived ones.

While implementing these changes faces political and technical challenges, including resistance from industry groups concerned about increased compliance costs, targeted pilot programs could demonstrate their feasibility and value. Focusing initially on high-impact areas like maternal health, AI would allow regulators to refine approaches before broader rollout, building evidence to counter objections. The alternative—allowing the current patchwork of voluntary standards to persist—risks perpetuating harms that undermine trust in medical AI.

Building Equitable AI Systems

The foundational issue plaguing women's health AI is the lack of robust, representative data. A multinational initiative could systematically address these gaps by funding the creation of open-access datasets for historically underrepresented conditions like endometriosis or perinatal mood disorders. These datasets must be curated with input from diverse populations, including marginalised communities often excluded from medical research. Parallel efforts should establish data cooperatives—patient-controlled repositories where women can securely share their health information on their terms, deciding how it's used and by whom. For rare conditions where large datasets are inherently scarce, advances in synthetic data generation could help, provided these techniques are rigorously validated to ensure they capture true biological variability rather than amplifying artefacts.

Incentivising inclusive AI development requires structural changes to current innovation ecosystems. Patent systems could offer extensions for algorithms demonstrably trained on diverse datasets, rewarding rather than penalising the extra effort needed for equitable development. Funding agencies should prioritise teams with gender-balanced leadership and community-embedded research practices, as homogeneous teams consistently overlook critical design considerations. Certification programs for bias-mitigated algorithms—developed through transparent, participatory processes—would help healthcare providers identify tools meeting higher equity standards, creating market pressure for improvement.

Resource constraints inevitably require phased approaches, making strategic prioritisation essential. Initial efforts should focus on creating open-access datasets for conditions where data gaps most severely impact outcomes, such as cardiovascular disease in women or racial disparities in maternal mortality. These foundational resources would enable cascading benefits across multiple research and development initiatives, maximising the impact of limited funding.

Ensuring Ethical Implementation

Effective governance structures must bridge the gap between ethical principles and daily clinical practice. This begins with mandatory independent ethics review boards for all clinical AI deployments, during development and throughout the product lifecycle. These boards should include technical experts, patient advocates, social scientists, and frontline clinicians who understand the real-world contexts where these tools will be used. Patient advocacy councils could assist in overseeing implementation, ensuring algorithms don't reinforce existing power imbalances in healthcare relationships. Accountability mechanisms must be established upfront, specifying recourse for individuals harmed by algorithmic decisions, whether through erroneous risk assessments or privacy breaches.

Capacity building represents the often-neglected cornerstone of ethical implementation. Healthcare workers need training beyond the technical operation of AI tools to encompass gender-sensitive interpretation of their outputs—recognising, for example, when an algorithm's recommendation might reflect historical underdiagnosis patterns rather than current clinical reality. Public education campaigns should demystify AI's role in women's health, empowering patients to ask informed questions about how algorithms influence their care. Community health workers—particularly in marginalised regions—require personalised support as they increasingly mediate between complex AI systems and populations with varying levels of health literacy.

Given resource limitations, prioritising frontline providers in underserved areas offers the biggest potential for equitable impact. These providers often serve as the sole healthcare access point for vulnerable populations, making their ability to engage with AI tools especially consequential. Training programs should be co-designed with these providers to address their specific challenges, such as intermittent connectivity or language barriers. By anchoring implementation ethics in the realities of those facing the steepest healthcare disparities, we can steer AI's trajectory toward genuine empowerment rather than high-tech exclusion.

Conclusion

AI’s integration into women's healthcare presents extraordinary opportunities and significant ethical challenges that demand careful navigation. While these technologies hold immense potential to improve health outcomes and reduce longstanding disparities, their benefits will not materialise automatically or equitably without intentional intervention. Achieving meaningful, ethical AI implementation requires deliberate and coordinated action across multiple fronts - from fundamentally transforming how health data is collected and utilised to prioritising establishing robust, gender-sensitive regulatory frameworks. Policymakers, healthcare leaders, and technology developers must engage in genuine co-creation processes to develop AI tools that demonstrably improve healthcare access and outcomes for women across diverse backgrounds and life stages. As AI systems become increasingly embedded in healthcare infrastructure, the governance structures we establish today will have lasting consequences, shaping health equity for future generations. By making inclusivity, transparency, and accountability foundational principles - rather than afterthoughts - we can harness AI's transformative potential to create a future where every woman, regardless of race, socioeconomic status, or geographic location, can access high-quality, personalised healthcare. However, the ambitious scope of these necessary changes underscores the reality that success will require unprecedented multilateral cooperation and phased implementation, with clear, measurable benchmarks to track progress amid inevitable resource constraints and competing priorities. This complex undertaking demands sustained commitment from all stakeholders to ensure that technological advancements translate into tangible improvements in women's health worldwide.

Bibliography

Adamson, A.S. and Smith, A. (2018) 'Machine learning and health care disparities in dermatology', JAMA Dermatology, 154(11), pp.1247-1248. https://doi.org/10.1001/jamadermatol.2018.2348

Cirillo, D., Catuara-Solarz, S., Morey, C., Guney, E., Subirats, L., Mellino, S., Gigante, A., Valencia, A., Rementeria, M.J., Chadha, A.S. and Mavridis, N. (2020) 'Sex and gender differences and biases in artificial intelligence for biomedicine and healthcare', npj Digital Medicine, 3(1), pp.1-11. https://www.nature.com/articles/s41746-020-0288-5

Leslie, D. (2019) Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector. London: The Alan Turing Institute. https://www.turing.ac.uk/sites/default/files/2019-06/understanding_artificial_intelligence_ethics_and_safety.pdf

Maher, N.A., Senders, J.T., Hulsbergen, A.F., Lamba, N., Parker, M., Onnela, J.P., Bredenoord, A.L., Smith, T.R. and Broekman, M.L. (2022) 'Passive data collection and use in healthcare: A systematic review of ethical issues', International Journal of Medical Informatics, 159, p.104642. https://pubmed.ncbi.nlm.nih.gov/31445262/

McKinney, S.M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., Back, T., Chesus, M., Corrado, G.C., Darzi, A. and Etemadi, M. (2020) 'International evaluation of an AI system for breast cancer screening', Nature, 577(7788), pp.89-94. https://pubmed.ncbi.nlm.nih.gov/33057216/

Obermeyer, Z., Powers, B., Vogeli, C. and Mullainathan, S. (2019) 'Dissecting racial bias in an algorithm used to manage the health of populations', Science, 366(6464), pp.447-453. https://www.science.org/doi/10.1126/science.aax2342

Sandall, J., Tribe, R.M., Avery, L., Mola, G., Visser, G.H., Homer, C.S., Gibbons, D., Kelly, N.M., Kennedy, H.P., Kidanto, H. and Taylor, P. (2018) 'Short-term and long-term effects of caesarean section on the health of women and children', The Lancet, 392(10155), pp.1349-1357. https://www.healthynewbornnetwork.org/hnn-content/uploads/Caesarean2.pdf

Topol, E. (2019). Deep medicine: how artificial intelligence can make healthcare human again. New York: Basic Books. https://dl.acm.org/doi/10.5555/3350442

Vayena, E., Blasimme, A. and Cohen, I.G. (2018) 'Machine learning in medicine: Addressing ethical challenges', PLoS Medicine, 15(11), e1002689. https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.1002689

Wahl, B., Cossy-Gantner, A., Germann, S. and Schwalbe, N.R. (2018) 'Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings?', BMJ Global Health, 3(4), e000798. https://gh.bmj.com/content/3/4/e000798

World Health Organization (2021) Ethics and governance of artificial intelligence for health. Geneva: WHO. https://www.who.int/publications/i/item/9789240029200