2026 Volume: 02 Issue: 08 www.ijrpa.com ISSN 2456-9995 Review Article

International Journal Research Publication Analysis

Page: 01-16

DIGITAL INEQUALITY AND SOCIAL STRATIFICATION IN THE AGE OF ARTIFICIAL INTELLIGENCE

***Malay Khan**

Editor-In-Chief, Star National Productions.

Article Received: 20 July 2026
Article Revised: 09 August 2026
Published on: 29 August 2026

***Corresponding Author: Malay Khan**
Editor-In-Chief, Star National Productions.

DOI: https://doi-doi.org/101555/ijrpa.4968

ABSTRACT

The rapid proliferation of artificial intelligence (AI) across economic, educational, and governance sectors has inaugurated a new phase in the history of technological inequality. While AI promises efficiency, objectivity, and democratized access to knowledge, its deployment within societies marked by deep-seated social hierarchies risks not merely replicating but actively amplifying existing inequalities. This paper examines the intersection of digital inequality and social stratification in the age of AI, with particular focus on India as a critical case study. Drawing on recent empirical research, national survey data, and theoretical frameworks from stratification economics and critical algorithm studies, the paper argues that AI is creating a new axis of social stratification—what may be termed the "intelligence divide"—that intersects with and deepens existing inequalities of caste, class, gender, and geography. The paper documents how algorithmic bias manifests along caste and gender lines, how AI-mediated education and employment systems reproduce structural disadvantages, and how the digital infrastructure that underpins AI access remains unevenly distributed across social groups. It further examines the emergence of a "data underclass" whose exclusion from AI training datasets perpetuates representational invisibility. The paper concludes by evaluating policy responses, including the IndiaAI Mission and digital public infrastructure initiatives, and argues that meaningful AI equity requires moving beyond techno-solutionist approaches toward structural reforms that address the social foundations of inequality.

KEYWORDS: Digital Inequality, Artificial Intelligence, Social Stratification, Caste, Gender, Digital Divide, Algorithmic Bias, India, Intelligence Divide, Data Justice.

1. INTRODUCTION

Artificial intelligence has become pervasive in shaping decision-making infrastructures across sectors—from education and employment to healthcare, policing, and financial services. AI systems are frequently celebrated for their promise of efficiency and neutrality, offering the possibility of objective, data-driven decisions that transcend human bias. Yet an expanding body of literature exposes that algorithmic systems perpetuate and even intensify social biases. This paradox—the promise of neutrality against the reality of reproduced inequality—lies at the heart of contemporary debates about digital inequality and social stratification.

The stakes of this inquiry are particularly high in India, a nation that has positioned itself as a global AI powerhouse while simultaneously grappling with some of the world's most entrenched social hierarchies. India is one of the largest markets for AI technologies, ranking among the top users of platforms like OpenAI's ChatGPT and Anthropic's Claude. The IndiaAI Mission, launched in 2024, aims to build a comprehensive ecosystem for AI innovation, treating AI not as a luxury but as a public good. Yet this technological expansion occurs within a socio-political landscape deeply shaped by caste, religion, and economic class.

The central argument of this paper is that AI is creating a new dimension of social stratification—an "intelligence divide"—that builds upon and deepens existing digital inequalities. This divide operates at multiple levels: access to AI infrastructure, exposure to AI-mediated opportunities, representation within AI training data, and vulnerability to algorithmic discrimination. For marginalized communities already disadvantaged by caste, gender, or geography, AI threatens to become not a tool of empowerment but a mechanism of exclusion.

This paper proceeds as follows. Section 2 reviews the theoretical literature on digital inequality and social stratification, establishing the conceptual framework for understanding AI as a stratifying force. Section 3 examines the existing digital divide in India, mapping the uneven distribution of access, skills, and usage across social groups. Section 4 analyzes how AI systems reproduce and amplify social inequalities through algorithmic bias, focusing on caste and gender. Section 5 explores the transformation of education and employment in the AI era, documenting how these sectors are becoming sites of intensified stratification. Section 6 examines the emergence of a "data underclass" and the politics of representational invisibility. Section 7 evaluates policy responses, including the IndiaAI Mission and digital public infrastructure initiatives. Section 8 concludes with implications for research and practice.

2. Theoretical Framework: Digital Inequality and Social Stratification

2.1 Conceptualizing Digital Inequality

Digital inequality is a multidimensional phenomenon that extends beyond mere access to technology. Following the work of scholars such as Van Dijk and Hargittai, digital inequality encompasses three interrelated dimensions: access (material resources), skills (digital literacy and competencies), and usage (meaningful participation and outcomes). This framework recognizes that connectivity alone does not ensure inclusion; what matters is how meaningfully individuals and communities participate in the digital world.

The National Sample Survey's Comprehensive Modular Survey on Telecom, 2025, provides a granular picture of these dimensions in India. While mobile phone usage among youth aged 15-29 is nearly universal at 97.1%, ownership tells a different story: only 73.4% of Indians own a mobile phone, with rural ownership (69.3%) lagging behind urban (82%) and female ownership (63%) lagging even further behind male ownership (83.3%). This ownership gap has serious implications for autonomy and privacy, often forcing women to rely on male family members to participate in education, access services, or make independent decisions online.

Skills inequality is even more pronounced. While 85.1% of youth can send emails or messages with attachments, only 32.2% have created an electronic presentation, and just 22.9% can draft documents using word-processing software. These figures suggest that digital engagement among Indian youth is largely social and transactional, not productive or generative. This skills gap has direct implications for AI readiness: individuals who lack foundational digital competencies are unlikely to harness AI's potential for cognitive enhancement or economic advancement.

2.2 Stratification Theory in the AI Era

Social stratification theory has long examined how societies organize inequality along axes of class, status, and power. In the digital age, scholars have extended these frameworks to understand how technology mediates stratification processes. The concept of the “digital divide” has evolved from a binary distinction between haves and have-nots to a recognition of layered inequalities in access, skills, and outcomes.

AI introduces a qualitatively new dimension to digital stratification. Unlike earlier information and communication technologies, AI is not merely a tool that users employ; it is an infrastructure that increasingly makes decisions about users. AI systems determine who is perceived as employable, credible, risky, or deserving. They allocate opportunities, assess performance, and predict behavior. When these systems are trained on data that reflects historical inequalities, they encode and automate those inequalities at scale.

Scholarship on AI and inequality has emerged mainly from Western academia, which inadequately captures the distinct social contexts of non-Western societies. In India, caste, gender, language, regional identity, and minority status constitute social situations fundamentally different from those in the West. Algorithmic decision-making does not function in isolation; it intersects with social hierarchies that structure access to resources, recognition, and rights. This intersectional perspective is essential for understanding how AI systems recalibrate social divisions in context-specific ways.

2.3 The Intelligence Divide

A growing body of research suggests that AI is transforming the traditional digital divide into a deeper "intelligence divide" between those who can harness cognitive enhancement and those who cannot. This concept captures the qualitative leap from access to information to the capacity to leverage AI for learning, problem-solving, and opportunity generation.

The intelligence divide operates through multiple mechanisms. First, access to AI tools remains highly uneven, with advanced AI applications requiring bandwidth, updated hardware, and cloud connectivity that urban, well-resourced institutions have solved but rural ones still lack. Second, AI literacy—the ability to understand, evaluate, and effectively use AI systems—is concentrated among already privileged populations. Third, AI systems themselves are optimized through interaction data; those who cannot engage regularly become "data invisible," and algorithms evolve primarily on patterns from already advantaged users.

The compounding effect is merciless: AI multiplies advantage for those already ahead while multiplying disadvantage for everyone else. This dynamic threatens to create a permanent stratification between those who are subjects of AI (actively shaping and benefiting from it) and those who are objects of AI (passively subjected to its decisions).

3. The Existing Digital Divide: Mapping Inequality in India

3.1 The Urban-Rural Divide

India's digital landscape is characterized by profound spatial inequalities. While national statistics suggest remarkable progress—over 90% of individuals aged 15-29 report using the internet—these aggregates mask deep disparities. Smartphone ownership stands at 82.1% in rural areas compared with 91.3% in urban areas. More significantly, only 14% of urban households have optical fibre connectivity compared with just 3.2% in rural areas.

The quality of internet connection matters as much as its availability. High-speed, reliable internet remains concentrated in cities, while rural areas grapple with slower, less stable connections. This infrastructure gap has direct implications for AI access, as most AI tools demand bandwidth and low-latency connectivity that rural India cannot provide.

A study examining the digital divide at the sub-national level found that the “Digital India” growth story is far from equitable, with low-income states and rural populations deserving greater attention. The researchers constructed Digital Infrastructure Indices and Digital Skills Indices for 18 Indian states, finding persistent and significant rural-urban gaps across both dimensions.

3.2 The Caste Divide

Perhaps the most distinctive feature of India's digital inequality is its caste dimension. Even when geography is held constant, internet access, its quality, and patterns of use differ significantly across caste groups. Upper-caste households lead India's digital landscape, with home internet access reaching 94.5% in urban areas and 86.4% in rural areas. At the other end of the spectrum, Scheduled Tribe (ST) households face the sharpest digital disadvantage: only 78.4% of rural ST homes have internet access, with the rural-urban gap standing at 9 percentage points—the widest among all caste groups.

Optical fibre connectivity reveals an even starker caste hierarchy. In urban India, fibre internet reaches 1 in 5 upper-caste households (20.4%), but only 1 in 8 ST households (12.8%). Upper-caste households are over 2.5 times more likely than SC and ST households to have fibre connections. Infrastructure gaps leave 8% of ST households offline—nearly four times the rate of upper-caste households (2.1%).

This caste-based digital inequality is not merely a matter of access; it shapes the very skills that determine economic opportunity. A study drawing on data from over 11 lakh individuals found that Scheduled Castes and Scheduled Tribes are at the most disadvantageous position in ICT skill acquisition. Persons who do not possess any ICT skills constitute 89.49% of STs, 86.62% of SCs, 81.73% of OBCs, and 73.71% of others. Rural residence and lower castes carry the largest burden of unequal access.

3.3 The Gender Divide

Gender constitutes another fundamental axis of digital inequality in India. Female mobile phone ownership (63%) lags behind male ownership (83.3%) by 20 percentage points. This ownership gap forces women to rely on male family members to access digital services, limiting their autonomy and privacy.

However, the gender gap in internet usage is narrowing, with young women's internet use surging from 77.1% in 2022-23 to 91.3% in 2025. This is not merely a technological trend but a social transformation. Yet usage patterns remain gendered: while 26% of male respondents use the internet primarily for entertainment, the proportion among females is notably higher at 36%. Women are less likely than men to use the internet for productive or informational purposes.

Skills inequality is even more pronounced. Across all skill levels, females are nearly 50% less likely than their male counterparts to have ICT competencies. Nationally, 22.78% of males have ICT skills compared to 13.91% of females. This gap has direct implications for AI participation: women are less exposed to the digital economy not because they are protected against job loss but because they are not present in the digital economy to begin with.

3.4 Intersectional Inequalities

The axes of inequality—rural-urban, caste, and gender—intersect to produce compounded disadvantages. A Scheduled Tribe woman in rural India faces barriers that are qualitatively different from and more severe than those faced by an upper-caste man in an urban area. Research on internet use among women in India found that education is the largest contributor to internet inequality (31.9%), followed by wealth (25.67%), place of residence (23.16%), and caste (1.10%). These findings underscore the need for intersectional analysis that captures how multiple axes of stratification interact.

4. Algorithmic Bias and the Reproduction of Inequality

4.1 Caste Bias in Large Language Models

The biases embedded in India's social structure do not remain external to AI systems; they are actively encoded and amplified within algorithmic architectures. Research on large language models (LLMs) in the context of caste and occupational identity in India has documented systematic biases that reflect and reinforce caste hierarchies.

A study by Zaveri and Shah (2025) examined three types of LLM biases through five studies covering a comprehensive set of occupations across all of India's districts. The findings are striking. First, individuals from marginalized caste groups are significantly under-represented in LLM output compared to their share in India's working population—a phenomenon that potentially reflects India's digital divide. Second, corrective measures to increase representation introduce other sources of errors, including association bias where marginalized castes are linked to occupations that require lower education levels and provide lower pay. Third, the models demonstrate selection bias with a higher probability of shortlisting resumes with names from dominant caste groups.

This last finding has profound implications for recruitment and hiring, where generative AI is becoming increasingly important as a cost-saving measure. If AI systems systematically favor upper-caste names in resume shortlisting, they will perpetuate and even amplify caste-based occupational segregation. An MIT Technology Review investigation similarly found that caste bias is rampant in OpenAI's products, with ChatGPT rendering Dalits as dirty, poor, and performing only menial jobs.

4.2 Gender and Intersectional Bias

Gender discrimination by AI is severe, particularly in patriarchal societal structures. Algorithmic discrimination is not confined to gender alone; it intersects with caste, religion, region, and language, creating situations in which populations become unevenly vulnerable. For example, AI hiring systems, used by 99% of Fortune 500 companies and increasingly by Indian firms, systematically filter out Muslim applicants. University of Washington research found that AI systems consistently ranked resumes with Muslim-associated names lower.

The intersectional nature of AI bias is further documented by research on STEM education content generated by LLMs. One study established that LLM-generated STEM content systematically disadvantages marginalized student profiles, with the gap between the most privileged and most marginalized profiles reaching 2.55 grade levels. This educational bias compounds existing inequalities in access to quality education and digital resources.

IIT Madras has released a dataset called IndiCASA (IndiBias-based Contextuality Aligned Stereotypes and Anti-stereotypes) comprising 2,575 human-validated sentences covering the demographics of caste, gender, religion, disability, and socio-economic status. The dataset addresses the inability of global models to capture the nuances of societal biases in the subcontinent. Evaluation of 14 popular LLMs on India-centric benchmarks reveals strong negative biases against marginalized identities, with models frequently reinforcing common stereotypes.

4.3 Structural Oppression and Algorithmic Governance

The bias in AI systems is not merely a technical problem; it reflects and reproduces structural oppression. As one systematic review argues, algorithmic bias arises from India's caste, class, gender, and religious hierarchies—not from code alone. Global AI governance models overlook these local inequalities that shape algorithmic outcomes. The exclusion of marginalized groups results from intertwined social, institutional, and technical flaws.

This structural perspective challenges the dominant techno-solutionist approach to AI governance. India's regulatory model prioritizes innovation and efficiency over rights, accountability, and inclusion. The analysis suggests that overcoming the crisis of legitimacy surrounding algorithmic governance requires abandoning the illusion of neutrality and advancing a data justice framework capable of monitoring and transforming the class consequences of algorithmic rule.

5. AI, Education, and the Future of Work

5.1 The Two-Tier Classroom

Education is perhaps the domain where the intelligence divide is most visibly manifest. As urban schools embrace artificial intelligence and digital learning, millions of children in rural India remain deprived of basic infrastructure, creating a new educational divide that could shape future economic inequalities.

In the glistening expanses of India's Tier-1 metros, students return to automated personalized learning modules, smart boards, and institutions piloting partnerships with global tech entities to integrate LLMs into daily pedagogy. But across the city limits in rural and Tier-3 India, the reopening of schools means a return to a struggle against basic amenities: cracked plaster, non-functional hardware, intermittent electricity, and a near-total lack of internet connectivity. This is not a unified national student body but a deep-seated educational disparity that threatens to permanently split India's demographic divide into two distinct economic castes.

The Wadhwani Foundation has documented how AI risks becoming the most powerful amplifier of India's oldest educational inequalities. Without deliberate inclusive design, the traditional digital divide is being transformed into a deeper “Intelligence Divide” between those who can harness cognitive enhancement and those who cannot. Student A in an urban, well-resourced school logs into an integrated platform where AI agents analyze past performance, prescribe personalized pathways, and generate adaptive quizzes. Student B in a rural school attends a class with sporadic electricity, shared smartphones, patchy or no internet, and manual record-keeping.

The infrastructure chasm is fundamental. Most AI tools demand bandwidth, updated hardware, and cloud connectivity that urban schools solved years ago but rural ones still lack. Smartphone penetration exists, but “one device per child” and always-online models remain economically and practically unfeasible for most rural families. The awareness abyss is equally critical: teachers often associate AI only with robots or science fiction, parents view it as “too advanced” or irrelevant for their children, and students first encounter the term through social-media reels rather than productive tools.

5.2 The Data Invisibility Trap

Perhaps the most insidious mechanism of AI-driven educational inequality is the data invisibility trap. AI improves through interaction data. Rural students who cannot engage regularly become “data invisible”. Algorithms and knowledge bases evolve primarily on urban, affluent patterns, introducing bias that further optimizes tools for cities and alienates villages. A permanent “data underclass” risks forming.

The consequences are compounding. Learning gaps that AI could close for lagging students will instead widen dramatically. As AI embeds in national assessments, competitions, skill development, and future jobs, today's exclusion becomes tomorrow's lockout. A phased rollout that starts with well-resourced pilots risks creating generational leapfrogging: by the time rural areas receive scaled solutions, urban pedagogy will already be several iterations ahead. India cannot afford a K-shaped or two-tier learning curve where geography and income predict cognitive opportunity.

5.3 Caste and AI Exposure in the Labour Market

The unequal distribution of AI exposure in education translates directly into labour market inequality. Research mapping occupational AI-exposure indices to India's redesigned Periodic Labour Force Survey documents a steep caste gradient among 83,000 employed graduates. Graduates from Scheduled Castes and Scheduled Tribes are 0.24-0.37 standard deviations less exposed to AI than upper-caste graduates within the same district. One in four SC graduates and one in three ST graduates work in farm or elementary occupations untouched by AI.

Because AI exposure commands a wage premium of up to 20%, generative AI stands to widen, not narrow, India's caste earnings gap. This finding challenges the optimistic narrative that technology will automatically democratize opportunity. Instead, it suggests that AI will reward those already positioned to benefit from it while leaving behind those excluded from its infrastructure and applications.

The mechanisms driving this gap are multiple. Graduates from marginalized communities are more likely to attend institutions with limited digital infrastructure, less likely to have developed AI-relevant skills, and more likely to be channeled into occupations with low AI exposure. The cycle is self-reinforcing: limited exposure leads to limited skills, which leads to limited opportunities, which leads to continued limited exposure.

6. The Data Underclass and Representational Injustice

6.1 The Politics of Data Representation

AI systems are only as good as the data on which they are trained. When training data systematically underrepresents or misrepresents marginalized communities, the resulting AI systems perpetuate these representational failures. This is not a neutral technical limitation but a form of representational injustice with concrete material consequences.

The underrepresentation of marginalized caste groups in LLM output, documented by Zaveri and Shah, reflects a deeper problem: these groups are less present in the digital world, and thus less present in the data that trains AI systems. GPT-4 responses consistently overrepresent culturally dominant groups far beyond their statistical representation, despite prompts intended to encourage representational diversity. This overrepresentation of dominant groups and underrepresentation of marginalized groups creates a feedback loop: AI systems learn to associate certain identities with certain roles, and these associations then shape how AI systems evaluate and categorize individuals.

Language exclusion compounds representational injustice. Indian children have expressed concern that their languages and daily life contexts are poorly represented in generative AI outputs. With India's linguistic diversity—22 scheduled languages and hundreds of regional languages—the dominance of English and a few major Indian languages in AI training data marginalizes speakers of other languages. This has practical consequences: AI systems that cannot understand or generate content in local languages cannot serve the populations that speak them.

6.2 The Data Labor Divide

The creation of AI systems depends on a vast, often invisible workforce of data laborers who label, annotate, and validate training data. In India, this work is concentrated among well-educated workers, many with STEM university education, who have failed to find appropriate jobs in their actual fields. These “ghost workers” perform the essential but devalued labor that makes AI possible.

The data labor divide mirrors broader patterns of social stratification. Data labeling work is often outsourced to lower-cost locations and performed by workers from marginalized communities. These workers have little control over the data they produce or the systems they help build. They are subjects of the AI economy but not participants in its benefits—a new form of digital labor exploitation that reproduces colonial patterns of extraction.

This raises fundamental questions about justice in the AI era. Who benefits from AI? Who bears its costs? Whose labor makes it possible? Whose data trains it? And who is excluded from its benefits? The answers to these questions reveal the deep entanglement of AI with existing structures of inequality.

6.3 The Catch-22 of Exclusion

The exclusion of marginalized communities from AI systems creates a catch-22. To be included, communities need access to digital infrastructure, skills, and representation in data. But the very systems that could facilitate inclusion are designed on data that excludes them. This is not a problem that can be solved by technical fixes alone; it requires structural changes in how AI systems are designed, governed, and deployed.

As one panel at the India AI Impact Summit 2026 warned, without ethical design, data sovereignty, and community participation, AI could deepen inequality instead of solving it. Digital public infrastructure (DPI) is racing ahead, powered by black-box AI that risks mass exclusion. The challenge is to ensure that the infrastructure of the AI age is built with, not for, the communities it is meant to serve.

7. Policy Responses and the Path Forward

7.1 The IndiaAI Mission

The Indian government has launched ambitious initiatives to democratize AI access. The IndiaAI Mission, launched in 2024, aims to build a comprehensive ecosystem for AI innovation, treating AI not as a luxury or proprietary tool but as a public good. The mission includes providing a pool of thousands of graphics processing units (GPUs), expanding access to high-performance compute resources, and integrating AI with digital public infrastructure.

A white paper from the Office of the Principal Scientific Adviser outlines key enablers for democratizing AI access, including expanding access to high-quality, representative datasets; providing affordable and reliable computing resources; and integrating AI with Digital Public Infrastructure. The paper suggests that India should work to integrate its own digital public infrastructure into AI systems and open the gates for small players to participate.

These initiatives represent a significant commitment to inclusive AI development. However, critics argue that India's techno-solutionist regulatory model prioritizes innovation and efficiency over rights, accountability, and inclusion. The question is whether the infrastructure being built will serve all Indians equally or whether it will replicate the inequalities of the digital age.

7.2 Digital Public Infrastructure and Inclusion

Digital public infrastructure (DPI) has emerged as a key strategy for expanding AI access. Platforms like Bhashini, the government-backed language translation platform, have been cited by the UNDP as a key example of how DPI can transform AI accessibility. India's AI for Bharat initiative, closely linked with Bhashini, aims to make AI accessible in Indian languages and across diverse contexts.

The UNDP has noted that India is well-positioned to tackle digital inequality and shape an inclusive transition in AI for the world. Yet the same report acknowledges that the region spans a 200-fold gap between its richest and poorest nations, with strong digital ecosystems in some countries and limited connectivity, skills, and infrastructure in others.

The challenge is to ensure that DPI does not become another mechanism of exclusion. As one analysis warns, digital public infrastructure powered by black-box AI risks mass exclusion.

Without accessibility, data sovereignty, and community participation, DPI could deepen inequality instead of solving it.

7.3 Beyond Techno-Solutionism

Addressing AI-driven inequality requires moving beyond techno-solutionism—the belief that technological fixes can solve social problems. The structural inequalities that shape AI outcomes—caste, class, gender, geography—cannot be solved by better algorithms alone. They require structural reforms that address the social foundations of inequality.

This means investing in digital infrastructure in underserved areas, but it also means addressing the educational, economic, and social inequalities that determine who can benefit from that infrastructure. It means ensuring that AI systems are designed with input from affected communities and that they are subject to meaningful accountability. It means recognizing that AI is not neutral and that its deployment must be guided by principles of justice, not just efficiency.

The UNDP's 2025 Human Development Report highlights the emergence of an AI divide—a new form of inequality layered over existing development gaps. Addressing this divide requires investments in infrastructure, skills, and global influence in AI. But it also requires confronting the deeper inequalities that make some populations more vulnerable to AI-driven exclusion than others.

8. CONCLUSION

This paper has examined the intersection of digital inequality and social stratification in the age of artificial intelligence, with particular focus on India. The analysis reveals a troubling pattern: AI systems, far from being neutral tools of progress, are actively reproducing and amplifying existing social inequalities along axes of caste, gender, class, and geography.

The evidence is clear. The digital divide in India is not merely a matter of access but of meaningful participation, with caste, gender, and geography determining who can harness digital technologies for productive purposes. AI systems exhibit systematic biases that reflect and reinforce these inequalities, from underrepresentation of marginalized castes in LLM outputs to discriminatory outcomes in hiring and education. The intelligence divide—the gap between those who can leverage AI for cognitive enhancement and those who cannot—is emerging as a new axis of stratification that compounds existing inequalities.

The implications are profound. As AI becomes embedded in the infrastructures of education, employment, governance, and commerce, the exclusion of marginalized communities from AI's benefits and their disproportionate exposure to its harms will shape the trajectory of social stratification for generations to come. The emergence of a data underclass, invisible to the algorithms that increasingly govern life, threatens to create a permanent divide between those who are subjects of AI and those who are objects of it.

Policy responses, including the IndiaAI Mission and digital public infrastructure initiatives, represent important steps toward inclusive AI development. Yet they risk replicating the techno-solutionist approach that has characterized much of digital development policy, prioritizing innovation and efficiency over rights and inclusion. Meaningful AI equity requires moving beyond technological fixes to address the structural inequalities that shape who benefits from AI and who is harmed by it.

Future research should examine the lived experiences of communities excluded from AI systems, the mechanisms through which algorithmic bias operates in context-specific ways, and the conditions under which AI can become a tool of empowerment rather than exclusion. It should also critically evaluate policy responses, assessing whether they are achieving their stated goals of inclusion or merely reproducing existing inequalities in new forms.

The age of artificial intelligence presents both unprecedented opportunities and unprecedented risks. Whether AI becomes a force for greater equality or greater stratification depends not on the technology itself but on the social, political, and economic choices that shape its development and deployment. These choices must be guided by a commitment to justice, accountability, and the recognition that all communities deserve not just access to AI but meaningful participation in shaping its future.

REFERENCES

1. *Zaveri, J., & Shah, A. (2025). Caste and Occupational Identity in Large Language Models. Working Paper No. 724, Indian Institute of Management Bangalore.*
2. *AI in Social Divides: Data Policies, Structural Oppression, and Decision Automation in India. (2025). IEEE Explore.*
3. *The Great Digital Divide: Two Indias meet at the school gates. (2026). Tehelka.*
4. *Sarang, G. (2025). AI in Education: Why rural and Tier-3 India risks being left behind. Wadhwani Foundation.*

5. *Peering into India's digital divide. (2025, July 9). The Hindu.*
6. *Unplugged and Unseen: The Caste Divide in Connected India. (2025). CEDA, Ashoka University.*
7. *Mapping the digital divide across caste and class shows why skilling is highly skewed in India. (2025, November 23). The Hindu.*
8. *IIT Madras releases dataset to detect AI biases in LLMs. (2025, October 8). The Hindu.*
9. *AI in a Class-Diverse India: Rights, Representation, and Regulation. (2026). Jindal Global University.*
10. *Structural oppression and AI: A systematic review of data policy frameworks in India. (2025). ScienceDirect.*
11. *The Privilege of Exposure: Caste and Generative AI in India's Graduate Labour Market. (2025). arXiv.*
12. *OpenAI is huge in India. Its models are steeped in caste bias. (2025, October 1). MIT Technology Review.*
13. *India AI Impact Summit 2026: Does Digital Public AI Infrastructure Exclude The Disadvantaged? (2026, February 18). Outlook India.*
14. *AI has learned caste — and caste bias. (2026, February 18). Indian Express.*
15. *DECASTE: Unveiling Caste Stereotypes in Large Language Models Through Multi-Dimensional Bias Analysis. (2025). IJCAI.*
16. *How Deep Is Representational Bias in LLMs? The Cases of Caste and Religion. (2025). AAAI.*
17. *India's deep-tech democracy for inclusive AI. (2025, November 15). The Hindu.*
18. *Govt Releases White Paper On Democratising Access To AI Infrastructure. (2025, December 30). Zee News.*
19. *UNDP report says India model can shape inclusive AI transition. (2025, December 2). Fortune India.*
20. *Rethinking AI: The lessons for India. (2025, June 13). Deccan Herald.*
21. *Rural–Urban Digital Divide: Evidence From Indian States. (2025). Arora, R. & Sapre, N. Wiley.*
22. *Singh et al. (2025). AI Based Framework for Bridging the Digital Divide and Promoting Equality in India. Semantic Scholar.*
23. *FairI Tales: Evaluation of Fairness in Indian Contexts with a Focus on Bias and Stereotypes. (2025). ACL.*

24. *Artificial intelligence and the reproduction of inequality: The role of algorithmic bias in social class imaginaries. (2025). Semantic Scholar.*
25. *The AI Gap: How Socioeconomic Status Affects Language Technology Interactions. (2025). ACL.*