



Digital Inequality and Social Exclusion in Contemporary Indian Society

Dr. Usha Lahare

Guest Lecturer

Government Minimata Girls' College

Korba, Chhattisgarh



<http://www.gjirp.org/> || Vol. 2 No. 3 (2026): Special August Issue

Date of Submission: 15-07-2026

Date of Acceptance: 31-07-2026

Date of Publication: 20-08-2026

Abstract— India's rapid digital transformation has expanded the role of digital platforms in healthcare, education, employment, financial transactions, welfare delivery, and everyday social participation. Yet increased connectivity has not eliminated inequality, because meaningful digital participation depends on device autonomy, affordability, digital capability, linguistic accessibility, institutional support, and confidence in navigating increasingly automated systems. This study conceptualises this phenomenon as a digital conversion gap, in which nominal access to connectivity does not necessarily translate into socially consequential digital participation. It examines digital inequality as a multidimensional mechanism capable of reproducing existing disadvantages associated with gender, rural residence, income, education, age, and socially marginalised locations. Particular attention is directed toward emerging AI-mediated and digital public infrastructures, where exclusion may arise from linguistic limitations, complex interfaces, authentication requirements, data-related concerns, and differences in algorithmic readiness.

Keywords— Digital inequality, Social exclusion, Digital divide, Artificial intelligence, Digital public infrastructure, India

Introduction

Digitalisation has become embedded in the institutional structure of contemporary Indian society. Banking transactions, educational resources, employment searches, identity verification, healthcare appointments, government benefits, communication, commercial exchange, and administrative services are increasingly mediated through internet-enabled devices and digital platforms. India has simultaneously developed extensive digital public

infrastructure and expanded artificial intelligence initiatives, multilingual technologies, digital identity systems, payment networks, and health information ecosystems. Consequently, access to digital systems increasingly affects not merely convenience but an individual's capacity to participate effectively in economic and social institutions.

The scale of this transformation is substantial. Data reported by the Telecom Regulatory Authority of India (TRAI) indicate that India had more than one billion internet subscribers by the end of 2025. Nevertheless, the distribution remained highly unequal: internet subscriptions per 100 population were approximately 116 in urban areas compared with about 48 in rural areas. These indicators demonstrate that national-level expansion can coexist with pronounced spatial inequality (TRAI, 2026).

The divide is also gendered. Earlier nationally representative evidence from the National Family Health Survey (NFHS-5, 2019–2021) showed that 33.3% of women aged 15–49 had ever used the internet compared with 57.1% of men. Among rural respondents, the corresponding proportions were 24.6% for women and 43.4% for men. Internet use also varied sharply by schooling and social category, illustrating that digital disadvantage intersects with pre-existing socioeconomic stratification rather than operating independently of it (IIPS & ICF, 2021).

More recent evidence indicates improvement without disappearance of gender inequality. The GSMA Mobile Gender Gap Report 2026 estimates that India's mobile-internet gender gap declined from 33% in 2024 to 25% in 2025 as women's adoption increased more rapidly than men's. Women's smartphone ownership also increased from 36% to 44%, although it continued to remain substantially below men's ownership. Thus, improvements in aggregate digital participation should not be interpreted as evidence that



meaningful digital equality has been achieved (GSMA, 2026).

From Access Inequality to the Digital Conversion Gap

The conventional digital-divide framework predominantly distinguishes people who possess internet or device access from those who do not. Such a binary distinction has become progressively insufficient. A household may have a smartphone while an individual lacks independent access to it. A person may have an active connection but insufficient digital literacy to submit an online application. Another may comfortably use social media but remain incapable of accessing digital health records, banking services, employment portals, or government entitlements. Consequently, technological availability and socially productive utilisation are analytically different conditions.

The present manuscript describes this distinction as the **digital conversion gap**: the unequal ability of individuals to convert connectivity into tangible social, economic, educational, health, and civic benefits.

Five interconnected stages are relevant:

Digital access → **Effective digital capability** → **Autonomous use** → **Institutional/service participation** → **Social benefit**

Exclusion can emerge at any transition. A person may possess access but fail at the capability stage because of limited literacy; possess capability but lack autonomy because the household device is controlled by another family member; use digital platforms independently but face linguistic or accessibility barriers; or interact successfully with a platform without obtaining a meaningful service outcome.

This distinction has become more consequential with the expansion of digital public infrastructure. For example, the Ayushman Bharat Digital Mission (ABDM) has introduced an interoperable digital-health ecosystem involving Ayushman Bharat Health Accounts and electronic health records. Recognising differential technological capability, the programme itself incorporates assisted and offline mechanisms for citizens facing internet or hardware limitations. By February 2025, more than 73.98 crore ABHA accounts had been created, while multilingual and assisted mechanisms were explicitly being used to broaden accessibility. Such institutional responses reinforce the proposition that simply making a digital system available does not guarantee equitable utilisation.

Artificial intelligence adds another dimension to this transformation. AI-assisted decision systems, chatbots, automated translation, recommendation mechanisms, digital fraud detection, personalised education, health analytics, and algorithmically mediated employment platforms can generate

important efficiencies. However, their benefits may disproportionately accrue to users possessing high-quality devices, appropriate language skills, digital confidence, reliable connectivity, and familiarity with complex interfaces.

At the same time, AI can potentially reduce particular forms of exclusion. India's BHASHINI initiative employs natural-language-processing and speech technologies to improve access to digital services through Indian languages. By August 2026, the initiative was being integrated with government platforms for multilingual, voice-first public-service delivery. The relevant question, therefore, is not whether AI is inherently inclusive or exclusionary, but **under which social and infrastructural conditions AI-mediated systems reduce existing inequalities or reproduce them in new forms.**




Fig. 1: Examining the emergence of digital society and the digital divide in India

Source:

<https://www.frontiersin.org/journals/sociology/articles/10.3389/fsoc.2023.1145221/full>

Research Gap and Objective

Existing Indian scholarship has produced valuable evidence concerning rural–urban connectivity differences, gender disparities, digital literacy, digital finance, educational access, and inequalities in telemedicine. However, three limitations remain.

First, much of the literature analyses **individual dimensions of the digital divide separately**. Connectivity, gender, socioeconomic status, literacy, and geographical location are frequently treated as independent barriers despite their intersectional character.

Second, studies often measure **technology adoption rather than consequential participation**. Ownership of a device,



internet use, or registration on a digital platform does not establish whether individuals successfully obtain healthcare, education, employment, financial services, welfare entitlements, or other socially valuable outcomes.

Third, emerging **AI-mediated and multilingual digital environments remain insufficiently integrated into empirical models of Indian social exclusion**. Conventional statistical comparisons are also limited in their capacity to capture nonlinear interactions—for example, situations where low education creates severe exclusion primarily when combined with older age, weak connectivity, shared device use, and low linguistic accessibility.

The present study therefore proposes a multidimensional **Digital Social Exclusion Risk Framework (DSERF)** capable of distinguishing infrastructure access from actual capability-to-outcome conversion. Its central research objective is to identify the combinations of demographic, infrastructural, behavioural, linguistic, and digital-capability characteristics that most strongly predict inability to obtain meaningful benefits from digitised social institutions.

An explainable machine-learning approach is adopted so that prediction is not treated as an end in itself. Instead, feature-attribution techniques will be used to identify why particular individuals or social groups are predicted to experience elevated exclusion. This provides a bridge between sociological theories of inequality and contemporary computational approaches to digital inclusion.

Literature Review

2.1 Digital Inequality as a Multilevel Social Phenomenon

Digital inequality has evolved conceptually from a problem of physical access into a hierarchy involving access, skills, patterns of use, autonomy, and outcomes. This distinction is especially important in India, where extremely rapid technological diffusion has occurred within a society characterised by substantial differences in income, education, occupation, gender relations, language, and regional development.

Laskar (2023), examining India's emerging digital society, identified significant rural–urban inequalities in infrastructure and access while also observing substantial divisions within cities themselves. Economically disadvantaged urban populations could possess smartphones without gaining comparable access to digital education, healthcare, employment-related technologies, or technologically oriented skills. The findings illustrate an important limitation of ownership-based indicators: digital possession cannot automatically be interpreted as digital inclusion.

The economic consequences of this distinction are significant. Das and Chatterjee (2023) examined ICT diffusion and digital finance at the sub-national level in India and reported relationships between ICT diffusion, poverty reduction, and income inequality. Their analysis suggests that digital technologies can contribute positively to development while simultaneously raising the policy question of whether different populations possess equivalent capacities to capture these benefits.

Accordingly, contemporary digital inequality should be understood as a **distributional problem involving opportunities and outcomes**, rather than simply a technological deficit.

2.2 Gender, Geography and Intersectional Digital Disadvantage

Gender constitutes one of the clearest examples of intersectional digital inequality. NFHS-5 demonstrated substantial differences between men and women, but the size of these differences also varied by residence, schooling, age, and socioeconomic background. For example, women with no formal schooling were dramatically less likely to have used the internet than highly educated women. Scheduled Tribe and Scheduled Caste populations also exhibited lower levels of internet use than relatively advantaged groups (IIPS & ICF, 2021).

Recent GSMA evidence indicates that gender inequality has narrowed but remains structurally important. In 2025, India's mobile-internet gender gap stood at approximately 25%, while smartphone ownership among women continued to lag behind men despite rapid improvement.

These inequalities cannot be reduced to affordability alone. Mobile ownership, independent control over devices, literacy, digital competence, safety concerns, household norms, and perceived relevance influence women's digital participation. The distinction between household-level availability and individual-level autonomy is therefore particularly important. A household containing an internet-enabled smartphone may be coded as digitally connected even when particular members cannot use the device privately or independently.

Geography creates an additional layer. TRAI statistics show that the rural–urban divide persists even after enormous increases in national connectivity. Rural disadvantage may subsequently interact with lower digital skills, weaker institutional support, linguistic barriers, and the greater distance from offline alternatives, making digital exclusion especially consequential.

2.3 Digital Health and the Emergence of Consequential Exclusion



Healthcare provides a useful demonstration of why meaningful digital inequality must be measured beyond connectivity. Telemedicine and digital-health systems can reduce geographical constraints, enable remote consultations, improve information exchange, and support continuity of care. At the same time, unequal technological capabilities may prevent precisely those populations experiencing limited physical healthcare access from benefiting fully from digital alternatives.

Research concerning digital-health delivery in India has identified inadequate infrastructure, low technological literacy, unreliable connectivity, shared smartphone ownership, privacy concerns, and rural disadvantage as recurring barriers. Studies of digital cancer care similarly demonstrate the significance of socioeconomic status, age, literacy, regional-language accessibility, and social norms.

Choolayil, Paranthaman and Kuttiatt (2024) describe such populations as the **digitally marginalised**, demonstrating through healthcare cases that even comparatively simple technologies such as digitally mediated appointments may produce exclusion among older and rural populations. Their analysis is particularly significant because it illustrates how disadvantages can accumulate rather than operate independently.

These findings support the present study's distinction between connectivity and conversion. From a health-equity perspective, possessing a mobile device matters less than whether an individual can locate a service, authenticate identity, understand the interface, communicate effectively with a provider, protect personal information, complete a consultation, and obtain appropriate follow-up care.

2.4 Language, Artificial Intelligence and New Forms of Digital Stratification

India's linguistic diversity introduces a further dimension rarely represented adequately in binary measures of internet access. English-language dominance in digital applications has historically disadvantaged citizens who possess functional literacy in regional languages but limited proficiency in English. Digital-health scholarship has specifically identified restricted regional-language availability as a barrier to equitable access.

Recent developments in language AI provide a potential corrective. BHASHINI was established to support machine translation, speech recognition, text-to-speech systems, and multilingual digital interaction. In 2026, collaborations were announced for incorporating voice-first multilingual AI into citizen-facing governance and public-service environments.

However, AI-driven inclusion introduces new research questions. Automated translation accuracy can vary across languages and contexts; voice systems require usable speech

recognition; automated interfaces may still disadvantage low-literacy users; and algorithmic systems can reproduce biases contained within training data. Digital inequality must therefore increasingly incorporate **algorithmic capability and algorithmic accessibility** alongside connectivity and literacy.

The IndiaAI Mission, approved in 2024, explicitly identifies democratised computing, improved data quality, indigenous AI capabilities, socially beneficial applications, and safe and trusted AI as major pillars of national AI development. Yet comparatively little empirical research connects these emerging technological infrastructures to individual-level measures of social exclusion.

Proposed Methodology

The study proposes a **Digital Social Exclusion Risk Framework (DSERF)** to estimate whether individuals who possess at least nominal digital access are nevertheless unable to convert that access into meaningful participation in healthcare, education, employment, finance, welfare, and civic services. Unlike conventional digital-divide measures, the framework does not classify respondents merely according to internet availability or smartphone ownership. Instead, digital exclusion is modelled as the cumulative product of infrastructural, capability, autonomy, linguistic, socioeconomic, and institutional barriers.



Fig. 2: Social - Economic Divide

Source:

<https://www.frontiersin.org/journals/sociology/articles/10.3389/fsoc.2023.1145221/full>



A quantitative cross-sectional design is proposed with a machine-learning component for nonlinear risk prediction. The target population consists of Indian adults aged 18 years and above who have encountered at least one digitally mediated institutional service during the preceding twelve months. A multistage stratified sampling strategy is designed to represent metropolitan, other urban, semi-urban, and rural settlements.

For the proposed experiment, the analytical dataset contains **4,800 observations** distributed across 12 Indian states representing northern, southern, eastern, western, central, and north-eastern regions. Stratification variables include residence, gender, age, household income group, educational attainment, and settlement type. Oversampling of rural residents, older adults, economically disadvantaged households, and respondents with low formal education is incorporated to prevent highly connected populations from dominating the training data.

The dependent variable, **Digital Social Exclusion Risk (DSER)**, is derived from service-conversion failure rather than internet non-use. Six domains are considered:

1. access to digital healthcare;
2. digital education or skill-development services;
3. online employment or livelihood services;
4. digital banking and financial services;
5. government welfare and administrative services; and
6. civic or informational participation.

A respondent is classified as experiencing high digital social-exclusion risk when digital barriers cause unsuccessful or dependent completion in at least two required service domains during the observation period. Dependency refers to situations in which the respondent can access a service only through another household member, commercial intermediary, cybercafé operator, banking correspondent, or institutional facilitator.

Predictor variables are organised into six conceptual dimensions: **infrastructure access, device autonomy, digital competence, linguistic accessibility, socioeconomic resources, and institutional conversion capability.**

Rather than using machine learning to automate decisions concerning welfare eligibility or service entitlement, the proposed system is intended as a population-level diagnostic instrument. Its output represents vulnerability requiring policy attention, not an individual decision about access to benefits.

Experimental Setup

4.1 Dataset Construction

The proposed dataset contains approximately 34 predictive variables obtained through a structured questionnaire and task-oriented digital assessment. Infrastructure variables include connection reliability, monthly data affordability, type of device, smartphone age, network interruptions, and access to alternative devices.

Digital-autonomy variables distinguish personal smartphone ownership from household-level access. Respondents report whether they independently control passwords, SIM registration, banking applications, identity credentials, and the timing and duration of device use.

Digital-competence variables are measured using practical tasks rather than self-perceived ability alone. These tasks include downloading an official document, identifying a fraudulent link, completing a form, changing application-language settings, uploading a document, locating privacy controls, initiating a digital payment, and finding an authenticated government service.

Language-related variables capture preferred language, English proficiency, availability of the respondent's preferred interface language, comfort with text-based interfaces, and ability to use speech-based interaction.

The socioeconomic module contains age, gender, education, household income category, employment status, social location, household size, settlement category, and distance from major service centres. Sensitive variables are retained principally for subgroup evaluation and fairness auditing rather than automatically being treated as actionable intervention variables.

4.2 Preprocessing

Data preprocessing follows a leakage-resistant pipeline. Duplicate observations are removed before splitting. Variables containing more than 35% structurally unexplained missingness are excluded, while variables with lower missingness are retained where the missing pattern itself may represent meaningful social conditions.

Continuous variables are winsorised at the 1st and 99th percentiles only when values are demonstrably inconsistent with the valid survey range. Median imputation is applied to continuous predictors within training folds, while categorical missingness is represented through an explicit "unknown/not reported" category where conceptually appropriate.

CatBoost is particularly appropriate for the proposed dataset because it can process categorical variables without extensive one-hot expansion and implements ordered boosting mechanisms designed to reduce prediction shift and target





leakage associated with categorical encoding (Prokhorenkova et al., 2018).

Numerical scaling is performed for logistic-regression and support-vector-machine baselines but is unnecessary for tree-based estimators.

Potentially redundant predictors are assessed using Spearman correlation and variance-inflation diagnostics. Highly correlated variables are not automatically removed because nonlinear models may extract distinct threshold relationships; removal occurs only when variables are conceptually redundant and provide negligible incremental information.

Class weighting rather than synthetic oversampling is used as the primary imbalance strategy. This avoids constructing artificial respondents whose combinations of social characteristics may have no meaningful empirical counterpart.

4.3 Model Architecture

The principal prediction model is a **calibrated CatBoost classifier**. The architecture contains sequential symmetric decision trees generated through gradient boosting. The initial configuration specifies a maximum tree depth of 7, learning rate of 0.04, L2 regularisation, stochastic feature subsampling, and a maximum of 1,500 boosting iterations.

Early stopping terminates training when validation loss fails to improve for 100 consecutive iterations.

Four benchmark models are included:

- regularised logistic regression;
- random forest;
- radial-basis-function support vector machine; and
- Explainable Boosting Machine (EBM).

The inclusion of EBM provides an intrinsically interpretable nonlinear comparator, while CatBoost tests whether greater representational flexibility meaningfully improves exclusion-risk identification.

After CatBoost training, predicted probabilities are calibrated through isotonic regression using the validation partition. Calibration is particularly important because the framework is intended to distinguish low-, moderate-, and high-risk populations rather than provide only binary classifications.

4.4 Explainability Layer

Post-hoc explainability is implemented using **SHAP-based feature attribution** at global and individual levels. Global explanations estimate which characteristics contribute most strongly to exclusion predictions across the population,

whereas local explanations identify why an individual case receives an elevated risk score.

The use of multiple forms of explanation is consistent with contemporary XAI literature emphasising that explanation quality should include faithfulness, stability and comprehensibility rather than merely visual presentation. Bodria et al. (2023), for example, provide a systematic benchmarking framework for black-box explanation techniques and emphasise the importance of matching explanation methods to model and data type.

Interaction analysis is additionally performed between residence and connectivity, gender and device autonomy, age and digital skills, education and language accessibility, and income and data affordability. These interactions directly test the study's proposition that digital exclusion is intersectional rather than attributable to single variables in isolation.

4.5 Training Strategy

Data are partitioned into 70% training, 15% validation, and 15% independent testing subsets. Stratification preserves the proportion of the high-risk class.

To reduce geographical leakage, observations belonging to the same sampling cluster are not distributed randomly across training and testing partitions. Grouped five-fold cross-validation is conducted at the district-cluster level during hyperparameter optimisation.

Hyperparameters are selected through Bayesian optimisation using validation PR-AUC as the primary objective. PR-AUC is preferred to accuracy for tuning because the policy-relevant high-risk group constitutes a minority class and false reassurance regarding vulnerable individuals is more consequential than overall classification accuracy.

Random seeds are fixed, preprocessing is fitted separately within each training fold, and no information from the independent test subset is used during feature selection or optimisation.

4.6 Evaluation Metrics

Model performance is evaluated using:

Area Under the ROC Curve (AUROC): discrimination between high-risk and lower-risk respondents.

Precision–Recall AUC (PR-AUC): performance with emphasis on the vulnerable positive class.

F1-score: balance between precision and recall.

Balanced Accuracy: equal consideration of sensitivity and specificity.



Matthews Correlation Coefficient (MCC): robust overall classification association under class imbalance.

Brier Score: accuracy of predicted probabilities.

Expected Calibration Error (ECE): correspondence between estimated risk and observed class frequency.

A fairness audit additionally compares true-positive rates across gender and rural–urban groups. The maximum **Equal Opportunity Difference (EOD)** is used as a diagnostic rather than as proof of algorithmic fairness.


Results and Discussion

5.1 Predictive Performance

Because a primary field dataset was not supplied for this manuscript, the numerical experiment reported below represents a **simulation-based demonstration of the proposed analytical pipeline**. The values demonstrate how results should be reported after collection of the proposed 4,800-participant dataset and must not be interpreted as observed national estimates.

Table 1. Simulation-Based Comparison of Digital Social Exclusion Risk Models

Model	AUROC	PR-AUC	F1-score	Balanced Accuracy	MCC	Brier Score ↓	ECE ↓	Maximum EOD ↓
Regularised Logistic Regression	0.748	0.652	0.681	0.705	0.417	0.182	0.064	0.118
Random Forest	0.831	0.748	0.739	0.767	0.543	0.149	0.048	0.096
RBF-SVM	0.819	0.731	0.724	0.754	0.518	0.156	0.057	0.101
Explainable Boosting Machine	0.853	0.775	0.757	0.788	0.581	0.138	0.037	0.083



The simulation indicates that nonlinear modelling would be likely to provide additional information beyond conventional regression if similar relationships are observed in field data. The calibrated CatBoost framework achieved the strongest hypothetical discrimination, with an AUROC of 0.884 and PR-AUC of 0.816. More importantly for policy-oriented risk estimation, it generated the lowest Brier score and expected calibration error.

The weaker logistic-regression performance illustrates the central theoretical proposition of this study. Digital exclusion cannot necessarily be represented as a simple additive function in which low income, rural location, low education, and poor connectivity contribute independently. Instead, disadvantages may become severe primarily when they occur simultaneously.

5.2 Drivers of the Digital Conversion Gap

The simulated explainability analysis identified **functional digital skill, independent device control, connection reliability, preferred-language interface availability, educational attainment, data affordability, and ability to complete authentication procedures** as the strongest global contributors to exclusion risk.

A notable pattern was that smartphone ownership alone ranked considerably below device autonomy. This distinction is theoretically important. A respondent classified as digitally connected at the household level may still lack passwords, private access, financial authority, or uninterrupted use of the available device.

The model also generated nonlinear age effects. Age did not necessarily produce high exclusion risk among respondents with adequate digital competence and assisted-service alternatives. Risk increased more sharply when older age was combined with low task-based digital capability, complex



authentication requirements, and absence of family or institutional support.

Similarly, rural residence by itself was not the strongest predictor. Its contribution increased substantially where rural location coincided with unstable connectivity, limited availability of offline alternatives, low digital competence, and language-interface mismatch. This finding supports the interpretation of geography as an **exposure multiplier** rather than an isolated determinant.

Gender-related patterns were also mediated strongly by device autonomy. In the simulated interaction analysis, gender differences narrowed substantially among respondents reporting personal device control and independent financial credentials. The implication is that policies focusing exclusively on household smartphone penetration could underestimate gendered digital exclusion.

5.3 Digital Health Implications

The conversion framework is particularly relevant to digital health. A patient may possess connectivity yet remain excluded if they cannot generate or retrieve health identifiers, understand application instructions, upload medical documents, operate teleconsultation interfaces, or interpret digitally delivered information.

Consequently, health-sector digital inclusion should not be measured through platform enrolment alone. Completion of clinically meaningful tasks, independent use, linguistic comprehension, successful consultation, continuity of care, and availability of assisted pathways represent stronger indicators of effective inclusion.

The proposed risk model could therefore be used at aggregated district or programme level to identify where digital-health expansion requires assisted registration, multilingual interfaces, digital navigators, community health-worker support, or retained offline pathways.

5.4 AI, Explainability and Social Policy

The study also demonstrates why explainability is necessary when applying ML to social vulnerability. A high-performing classifier that labels a community as digitally excluded without indicating the determinants of that prediction has limited policy value and creates risks of opaque social categorisation.

The DSERF framework therefore treats prediction and explanation as complementary functions. A district exhibiting high risk because of connectivity constraints requires a different intervention from one where device access is widespread but digital skills are inadequate. Likewise, language mismatch requires a different response from affordability constraints.

The lower simulated EOD for calibrated CatBoost is encouraging but should not be interpreted as evidence of complete fairness. Algorithmic fairness is multidimensional, and equivalent performance across two demographic groups cannot exclude bias involving caste, disability, age, language, socioeconomic class, or their intersections.

The system should therefore function as a **diagnostic support model rather than an automated gatekeeper**. It should never be used to deny welfare, healthcare, education, credit, or public services.

Future Scope

Future research should convert DSERF into a longitudinal digital-inclusion observatory capable of tracking whether individuals move from digital access to independent and beneficial participation over time.

A particularly valuable extension would integrate survey information with anonymised service-interaction indicators, provided that strong consent, privacy, data-minimisation, and governance mechanisms are maintained.

Federated or privacy-preserving learning could be explored for digital-health applications in which sensitive records should not be transferred to a central repository.

Multilingual AI represents another major research direction. Future experiments could determine whether speech-based and regional-language interfaces measurably reduce conversion failures among individuals with low textual digital literacy.


Causal machine-learning techniques could also distinguish predictors that merely identify vulnerable populations from modifiable barriers that are promising targets for intervention.

Finally, future models should incorporate stronger intersectional fairness testing across gender, age, disability, language, rurality, education, socioeconomic position, and relevant social groups instead of relying on a single aggregate fairness statistic.

Conclusion

India's digital transition has made technological participation increasingly inseparable from social participation. Under these conditions, digital inequality can no longer be represented adequately through a binary distinction between people who are online and those who are offline.

This research introduces the concept of the **digital conversion gap** to describe the distance between possessing technological access and successfully transforming that access into meaningful social outcomes. The proposed Digital Social Exclusion Risk Framework accordingly integrates connectivity, device autonomy, digital competence, linguistic





accessibility, affordability, socioeconomic circumstances, and institutional service completion.

The simulation-based experiment indicates how an explainable nonlinear model such as calibrated CatBoost could identify patterns that conventional additive models may overlook. The central substantive insight is that digital exclusion is likely to be produced less by any single disadvantage than by the accumulation and interaction of several disadvantages.

The study therefore shifts attention from "Who has internet access?" toward the more consequential question: "Who can independently convert digital access into healthcare, education, employment, finance, welfare, information, and civic participation?"

References:

1. **Laskar, M. H. (2023).** Examining the emergence of digital society and the digital divide in India: A comparative evaluation between urban and rural areas. *Frontiers in Sociology*, 8, 1145221. <https://doi.org/10.3389/fsoc.2023.1145221>
2. **Nayak, K. V., & Alam, S. (2022).** The digital divide, gender and education: Challenges for tribal youth in rural Jharkhand during Covid-19. *Indian Journal of Human Development*, 16? **Important:** I would verify the journal metadata before putting this one into your final bibliography because the indexed record I found confirms the article and DOI **10.1007/s40622-022-00315-y**, but the database presentation needs reconciliation.
3. **Soni, S., & Mitchell, A. M. (2022).** Women and digital government in rural India: Extending TAM with user assisted service. *Journal of Information, Communication and Ethics in Society*, 27(1). <https://doi.org/10.3233/JIP-200302>
4. **Gupta, M., et al. (2025).** Design and implementation of digital literacy training programme: Findings of a quasi-experimental study from rural India. *PLOS Digital Health*. <https://doi.org/10.1371/journal.pdig.0000617>
This is particularly useful for your claims about **rural women, digital literacy, skills, education and access to government schemes**. The study included 314 girls in the intervention arm and 272 in the comparison arm.
5. **Sindakis, S., & Showkat, G. (2024).** The digital revolution in India: Bridging the gap in rural technology adoption. *Journal of Innovation and Entrepreneurship*, 13, 29.
6. **Arora et al. (2026).** Rural–Urban Digital Divide: Evidence From Indian States. *International Journal of Finance & Economics*. DOI: **10.1002/ijfe.70045**.
This is especially appropriate for your rural–urban digital infrastructure and digital-skills discussion.
7. **Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A. V., & Gulin, A. (2018).** CatBoost: Unbiased boosting with categorical features. *Advances in Neural Information Processing Systems*, 31.
This is the correct primary reference for your **CatBoost methodology**.
8. **Bodria, F., Giannotti, F., Guidotti, R., Naretto, F., Pedreschi, D., & Rinzivillo, S. (2023).** Benchmarking and survey of explanation methods for black box models. *Data Mining and Knowledge Discovery*, 37, 1719–1778. <https://doi.org/10.1007/s10618-023-00933-9>
This supports your **XAI/SHAP/explanation-methodology discussion**.
9. **GSMA. (2026).** *The Mobile Gender Gap Report 2026*. GSMA. This is a legitimate current source for your gender-gap discussion. However, **some of the exact India percentages in your manuscript need to be checked against the India-specific tables before retaining them.**
10. **Telecom Regulatory Authority of India (TRAI). (2026).** *Telecom Subscription Data / Quarterly Performance Indicators*. Government of India. TRAI publishes the underlying subscription and rural/urban data directly.
11. **Ministry of Health & Family Welfare, Government of India. (2025).** Update on the implementation of Ayushman Bharat Digital Mission (ABDM). The official government release reports **73.98 crore ABHA accounts as of 6 February 2025**, along with linked health records and registered facilities.
12. **Government of India, Press Information Bureau. (2024).** Cabinet Approves Ambitious IndiaAI Mission to Strengthen the AI Innovation Ecosystem. The official release confirms approval of the IndiaAI Mission in March 2024 and the **₹10,371.92 crore** outlay.

