
Fairness Considered Harmful: On the Non-portability of Fair-ML in India

Nithya Sambasivan, Erin Arnesen, Ben Hutchinson, Vinodkumar Prabhakaran
Google Research
USA
(nithyasamba, erinarnesen, benhutch, vinodkpg)@google.com

Abstract

Conventional algorithmic fairness is Western in its sub-groups, values, and optimizations. In this paper, we ask how portable the assumptions of this largely Western take on algorithmic fairness are to a different geo-cultural context such as India. Based on 36 expert interviews with Indian scholars, and an analysis of emerging algorithmic deployments in India, we identify three clusters of challenges that engulf the large distance between machine learning models and oppressed communities in India. We argue that a mere translation of technical fairness work to Indian subgroups may serve only as a window dressing, and instead, call for a collective re-imagining of Fair-ML, by re-contextualising data and models, empowering oppressed communities, and more importantly, enabling ecosystems.

1 Introduction

Fairness research in Machine Learning (ML) has seen rapid growth in recent years; however, it remains largely rooted in Western concerns and histories: the injustices they focus on (*e.g.*, along race and gender), the datasets they study (*e.g.*, ImageNet), the measurement scales they use (*e.g.*, Fitzpatrick), and the legal tenets they draw from (*e.g.*, equal opportunity). This implicitly Western take is troublingly becoming a *universal* ethical framework for ML, *e.g.*, AI strategies from India [1], Tunisia [51], and Mexico [32] all derive from this work, but fails to account for the several assumptions conventional algorithmic fairness makes about the availability and efficacy of the surrounding institutions and infrastructures. However, these infrastructures, values, and legal systems cannot be naively generalised to various non-Western countries.

Let us consider the example of facial recognition technology, where demonstration of fairness failures resulted in bans and moratoria in the US. Several factors led to this outcome: decades of empiricism on proxies and metrics that correspond to subgroups in the West [21]; public datasets, APIs, and laws enabling analysis of model outcomes [11, 28]; an ML research/industry responsive to bias reports from users and civil society [8, 17]; the existence of government representatives glued into technology policy [53]; and an active media that scrutinizes downstream impacts of AI [28]. However, due to various cultural, ethnic, and infrastructural differences, these factors are often absent or irrelevant in much else of the world. Could fairness have structurally different meanings or mechanisms in non-Western contexts? How do social, economic, and infrastructural factors influence implementation of meaningful fairness?

In this position paper, we present insights from qualitative interviews with 36 Indian scholars and activists, and a discourse analysis of emerging algorithmic deployments in India. India is home to 1.38 billion people and their multiple languages, religions, cultural systems, and ethnicities. AI deployments are prolific in the public sector, *e.g.*, in predictive policing [13], facial recognition [19], and agriculture [34]. Despite this forward momentum, there is a dire lack of conversations on

advancing algorithmic fairness for such a large population. Our paper presents a first step towards filling this important gap. We contend that conventional Fair-ML approach may be inappropriate and inimical in India, if it does not engage with local structures. For a country as deeply plural, complex, and contradictory as India—where the distance between models and oppressed communities is large—optimising model fairness alone can be mere tokenism. We call for action along three critical pathways towards algorithmic fairness in India: *Recontextualising*, *Empowering*, and *Enabling* (see Figure 2). The considerations we present involve collective responsibility of inter-disciplinary Fair-ML researchers, and push the bounds of what is considered to be fairness research.

2 Methodology

We perform a synthesis of (i) 36 interviews with researchers and activists across various disciplines working with marginalized Indian communities, and (ii) observations of current algorithmic deployments and policies in India. We chose our respondents from diverse disciplines: Computer Science (11), Activism (9), Law and Public Policy (6), Science and Technology Studies (5), Development Economics (2), Sociology (2), and Journalism (1). All respondents had 10-30 years of experience working with on social justice in India, on areas including caste, gender, labour, disability, surveillance, health, and constitutional rights. In conjunction with qualitative interviews, we analyzed various algorithmic deployments and emerging policies in India, starting from 2009. We identified and analysed various Indian news publications, policy documents, grassroots fora, and prior research on automation in India. The systematic process followed for the collection, categorization, and synthesis of data to develop the key insights below are described in detail in Appendix A.

3 Summary of Findings

Our study finds that the conventional fairness approach may be inappropriate, insufficient, or even inimical in India, if it does not engage with local structures. In India where the distance between ML models and dis-empowered communities whom they aim to serve is large—via technical distance, social distance, ethical distance, temporal distance, and physical distance—a myopic take on localising ‘fair’ model outputs alone can backfire. We identify three high-level themes/clusters of challenges (Figure 1) that engulf the large distance between models deployed in India and the communities they impact. We summarize them below.

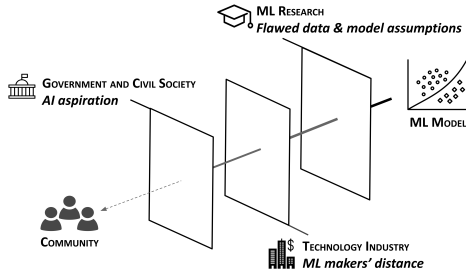


Figure 1: Challenges to algorithmic fairness in India: A Model-Community distance perspective

Flawed data and model assumptions: We found numerous ways data and model assumptions fail in India, primarily: *data distortions* owing to infrastructural challenges and technology usage patterns (e.g., SIM card sharing among women [47]) that makes datasets not faithful representations of people and phenomena; *data incompleteness* due to the models gearing for the ‘good’ data profiles that largely favor middle-class men—in a society where 50% do not have Internet, especially women and lower caste people; *axes of discrimination* and their proxies and indicators that differ from how they manifest in the West (for a detailed account, see Appendix 1); and *lack of Indic justice approaches* like reservations/quotas that present new challenges to fairness evaluations and interventions.

ML makers’ distance: Several respondents described a transactional mindset towards Indians, seeing them as agency-less data subjects that generated large-scale behavioural traces to improve ML models (also see [30]), resulting in poor recourse and redress for Indian users, especially for those dis-empowered by caste, religion, or gender. The human infrastructures who played a crucial role in providing recourse to marginalised Indian communities (e.g., street-side bureaucrats, and front line workers) were removed in ML systems. Many respondents also pointed out the lack of diverse representation in Indian tech workforce that the distant ML makers often fail to recognize.

AI aspiration: AI is aspirationally viewed in India, with high reliance for high-stakes domains, accompanied with high trust in automation, limited transparency and the lack of an empowered ecosystem to question algorithmic interventions. For instance, respondents pointed to the lack of a buffer zone of journalists, activists, and researchers to keep ML system builders accountable. Lack of

transparency and stakeholder participation, along with the its ‘neutral’ and ‘human-free’ associations lent misplaced credence to its algorithmic authority, making it further inscrutable.

4 Towards Algorithmic Fairness in India

To account for the challenges outlined above, we need to understand and design for *end-to-end* chains of algorithmic power, including how AI systems are conceived, produced, utilised, appropriated, assessed, and contested in India. To this end, we propose a research agenda where we call for action along three critical and contingent pathways towards successful AI fairness in India: *Recontextualising*, *Empowering*, and *Enabling* (see Figure 2). These pathways present new sociotechnical research challenges and require cross-disciplinary and cross-institutional collaborations.

4.1 Recontextualising data and models

Data plays a critical role in measurements and mitigation of algorithmic bias. We must be (even more than usual) sceptical of Indian datasets (due to challenges outlined above), and study how fairness research could handle the known distortions in the data as well as how we must account for data voids [23] for which statistical extrapolations might be invalid. The vibrant role played by human infrastructures in providing, negotiating, collecting, and stewarding data points should encourage us to re-imagine data as a dialogue rather than operations, i.e, products of both the beholder and the beheld. How might data consent work in this context? One approach could be to create transitive informed consent, built upon personal relationships and trust in data workers, with transparency on potential downstream applications. Ideas like collective consent [46], data trusts [39], and data co-ops may enhance community agency in datasets, while simultaneously increasing data reliability. Finally, we must question the categories and constructs we model in datasets, and how we measure them. When categories are appropriate for endemic goals (*e.g.*, caste membership for quotas), what form should their distributions and annotations take? Linguistically and culturally pluralistic communities should be given voices in these negotiations in ways that respect Indian norms of representation.

Model (un)fairness detection and mitigation should incorporate the prominent axes of historical injustices in India (see Appendix 1) and tackle the challenges in operationalising them for testing AI; *e.g.*, representational biases of caste and other sub-groups in NLP models, biases in Indic language NLP including challenges from code-mixing, Indian subgroup biases in computer vision, tackling online misinformation, benchmarking using Indic datasets, and fair allocation models in public welfare. It is important to note that operationalising fairness approaches from the West to these axes is often nontrivial. For instance, personal names act as a signifier for various socio-demographic attributes in India, however there are no large datasets of Indian names (like the US Census data, or the SSA data) that are readily available for fairness evaluations. Another important consideration is how the algorithmic fairness interventions work with the existing infrastructures in India that surrounds decision making processes; *e.g.*, in the context of restorative justice initiatives such as reservations/quotas? The quota system gives rise to the problem of matching under distributional constraints [27, 24, 12] that has not received much attention within the ML research.

4.2 Empowering communities

Participatory Fair-ML knowledge systems are critical to effectively fill the distance between ML makers and communities. Context-free assumptions in fairness research, whether in homegrown or international ML systems, can not just fail, but produce harm inadvertently when applied to different infrastructural and cultural contexts. Systems should be created that enable marginalised communities to participate in the production of knowledge systems about themselves in ML frameworks. Grassroots efforts like Deep Learning Indaba [3] and Khipu [4] are helpful examples in bootstrapping AI research in communities. Initiatives like Design Beku [2] and SEWA [6] are good decolonial examples of participatorily co-designing with under-served communities. Diverse values in algorithmic fairness can come from justice systems of indigenous Adivasis [54], *dharma* (social ethic) and *karma* (action) from Hinduism (note that Hindu concepts are sometimes rejected by Dalits for reinforcing hierarchy [40]) or (*sanghas*) (communism) espoused by Dr. B.R. Ambedkar [9], for example.

Low-resource considerations pertaining to India should guide ML researchers to think beyond making model outputs fair, and about India’s heterogeneous literacies, economics, and infrastructures that interfere with access to those fair outputs. Half the population of India is not online. Layers of the stack like interfaces, devices, connectivity, languages, and costs are all important in ensuring

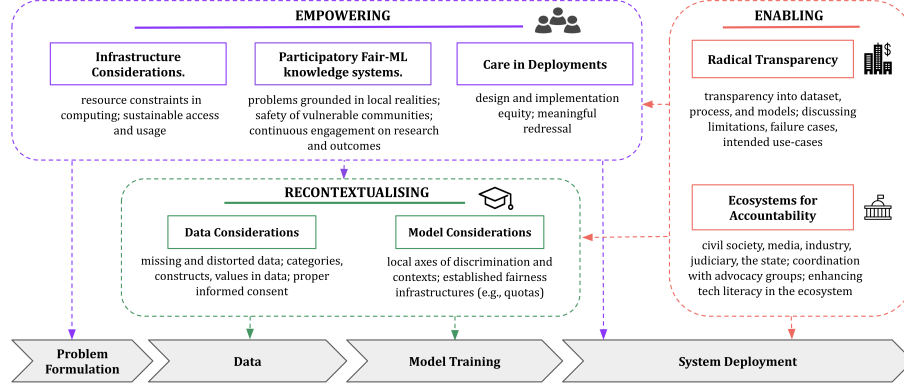


Figure 2: Research pathways for Algorithmic Fairness in India.

sustainable access and usage. Tackling resource constraints in computing disciplines like ICTD [50] and HCI4D [20] can help—where constraints have been embraced as design material, *e.g.*, delay-tolerant connectivity, low cost devices, text-free interfaces, partnership with civil society, and capacity building [16, 25, 29, 33, 48]. Data infrastructures to build, share and sustain localised datasets would also enhance access equity to holding AI accountable (*e.g.*, [5]).

Care in deployments should be a primary concern. Critiques were raised in our study on how neo-liberal ML followed a broader pattern of extraction from the ‘bottom billion’ data subjects and AI labourers. Low costs, large and diverse populations, and policy infirmities have been cited as reasons for following double standards in India, *e.g.*, in non-consensual human trials and waste dumping [31] (also see [37]). Past disasters in India, like the fatal Union Carbide gas leak in 1984—one of the world’s worst industrial accidents—point to faulty design and low quality standards for the ‘third world’ [52]. Similarly, overgrowth of ML capital, with unequal standards, inadequate safeguards, and dubious applications can lead to catastrophic effects (similar analogies have been made for content moderation [44, 48]). ML researchers should study how to ensure meaningful recourse within the ecosystems they are embedding, and respond and engage with Indian realities and user feedback.

4.3 Enabling fair ML ecosystems

Ecosystems for accountability requires enabling the civil society, media, industry, judiciary, and the state to meaningfully engage in holding AI interventions responsible. Moving from ivory tower research approaches to solidarity with various stakeholders through project partnerships, involvement in evidence-based policy, and policy maker education can help create a sustainable Fair-ML ecosystem based on sound empirical and ethical basis. Investigative journalism on algorithms is lacking in India, but sustainable partnerships may be created with rigorous media sources. Algorithmic fairness and ethics are not mainstream research topics in Indian academia today, but the academy has a pivotal role in advancing fairness in India. A concerted effort is required to create APIs, documentation, and socio-economic datasets to enable meaningful and equitable accountability in the ecosystem.

Radical transparency is required to counteract the inscrutability and authority created by the enthusiastic AI aspiration. Besides the role played by ecosystem-wide regulations and standards, radical transparency as a design principle should be espoused and enacted by Fair-ML researchers committed to India. Transparency on datasets, processes, and models (*e.g.*, [36, 22, 14]), openly discussing limitations, failure cases, and intended use-cases can help bolster the practical limits to applying computing to human problems. It is also incumbent upon those in power to discuss current Fair-ML interventions as a flawed and evolving approach. Such approaches can help move from the ‘magic pill’ role of fairness as a checklist for ethical issues in India, to a more pragmatic function.

5 Conclusion

We are currently facing a watershed moment for AI, with extravagant hype and pervasive applications, with nations being turned into AI superpowers and rulers [45]. We bring to light how efforts in the West to make AI fairer may not transfer readily elsewhere, specifically in the case of India, through an empirical bottom-up analysis. We present a roadmap to go beyond localising data and model

fairness alone, and present a holistic research agenda to operationalize Fair-ML in India. While it is important to incorporate Indian concepts of fairness into ML that impacts Indian communities, the broader call is to develop general approaches to fairness that reflect local considerations.

References

- [1] *National Strategy for Artificial Intelligence #AI4ALL*. Niti Aayog, 2018.
- [2] *Design Beku*, 2020. URL <https://designbeku.in/>.
- [3] *Deep Learning Indaba*, 2020. URL <https://deeplearningindaba.com/2020/>.
- [4] *Khipu AI*, 2020. URL <https://github.com/khipu-ai>.
- [5] *Lacuna Fund*, 2020. URL <https://lacunafund.org/>.
- [6] *SEWA*, 2020. URL <http://www.sewa.org/>.
- [7] Delna Abraham and Ojaswi Rao. 84% dead in cow-related violence since 2010 are muslim; 97% attacks after 2014 | indiaspend. <https://archive.indiaspend.com/cover-story/86-dead-in-cow-related-violence-since-2010-are-muslim-97-attacks-after-2014-2014>. (Accessed on 08/16/2020).
- [8] Amazon. We are implementing a one-year moratorium on police use of rekognition. <https://blog.aboutamazon.com/policy/we-are-implementing-a-one-year-moratorium-on-police-use-of-rekognition>, June 2020. (Accessed on 08/29/2020).
- [9] Bhimrao Ramji Ambedkar. *The Buddha and his dhamma: a critical edition*. Oxford University Press, 2011.
- [10] Bhimrao Ramji Ambedkar. *Annihilation of caste: The annotated critical edition*. Verso Books, 2014.
- [11] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias — propublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>, May 2016. (Accessed on 07/30/2020).
- [12] Itai Ashlagi, Amin Saberi, and Ali Shamel. Assignment mechanisms under distributional constraints. *Operations Research*, 68(2):467–479, 2020.
- [13] Abhishek Baxi. Law enforcement agencies in india are using artificial intelligence to nab criminals. <https://www.forbes.com/sites/baxiabhishek/2018/09/28/law-enforcement-agencies-in-india-are-using-artificial-intelligence-to-nab-criminals-heres-how>, September 2018. (Accessed on 08/30/2020).
- [14] Emily M Bender and Batya Friedman. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6: 587–604, 2018.
- [15] Alexander Bogner, Beate Littig, and Wolfgang Menz. *Interviewing experts*. Springer, 2009.
- [16] Eric Brewer, Michael Demmer, Bowei Du, Melissa Ho, Matthew Kam, Sergiu Nedevschi, Joyojeet Pal, Rabin Patra, Sonesh Surana, and Kevin Fall. The case for technology in developing regions. *Computer*, 38(6):25–38, 2005.
- [17] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, pages 77–91, 2018.
- [18] Ministry of Statistics Central Statistics Office and Programme Implementation. Women and men in india: A statistical compilation of gender related indicators in india. Technical report, Government of India, 2018.
- [19] Pranav Dixit. India is creating a national facial recognition system. <https://www.buzzfeednews.com/article/pranavdixit/india-is-creating-a-national-facial-recognition-system-and>, October 2019. (Accessed on 08/30/2020).
- [20] Susan Dray, Ann Light, A Dearden, Vanessa Evers, Melissa Densmore, D Ramachandran, M Kam, G Marsden, N Sambasivan, T Smyth, et al. Human–computer interaction for development: Changing human–computer interaction to change the world. In *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies, and Emerging Applications, Third Edition*, pages 1369–1394. CRC press, 2012.

- [21] Thomas B Fitzpatrick. The validity and practicality of sun-reactive skin types i through vi. *Archives of dermatology*, 124(6):869–871, 1988.
- [22] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *arXiv preprint arXiv:1803.09010*, 2018.
- [23] Michael Golebiewski and Danah Boyd. Data voids: Where missing data can easily be exploited. *Data & Society*, 2019.
- [24] Masahiro Goto, Fuhito Kojima, Ryoji Kurata, Akihisa Tamura, and Makoto Yokoo. Designing matching mechanisms under general distributional constraints. *American Economic Journal: Microeconomics*, 9(2): 226–62, 2017.
- [25] Kurtis Heimerl, Shaddi Hasan, Kashif Ali, Eric Brewer, and Tapan Parikh. Local, sustainable, small-scale cellular networks. In *Proceedings of the Sixth International Conference on Information and Communication Technologies and Development: Full Papers-Volume 1*, pages 2–12, 2013.
- [26] IDSIN. Two thirds of india’s dalits are poor - international dalit solidarity network. <https://idsn.org/two-thirds-of-indias-dalits-are-poor/>, July 2010. (Accessed on 08/13/2020).
- [27] Yuichiro Kamada and Fuhito Kojima. Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, 105(1):67–99, 2015.
- [28] Ava Kofman. How facial recognition can ruin your life – intercept. <https://theintercept.com/2016/10/13/how-a-facial-recognition-mismatch-can-ruin-your-life/>, October 2016. (Accessed on 07/30/2020).
- [29] Neha Kumar and Richard J Anderson. Mobile phones for maternal health in rural india. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 427–436, 2015.
- [30] Michael Kwet. Digital colonialism is threatening the global south | science & technology | al jazeera. <https://www.aljazeera.com/indepth/opinion/digital-colonialism-threatening-global-south-190129140828809.html>. (Accessed on 08/29/2020).
- [31] Ruth Macklin. *Double standards in medical research in developing countries*, volume 2. Cambridge University Press, 2004.
- [32] Emma Martinho-Truswell, Hannah. Miller, Isak Nti Asare, Andre Petheram, Richard (Oxford Insights) Stirling, Constanza Gómez Mont, and Cristina (C Minds) Martinez. Towards an ai strategy in mexico: Harnessing the ai revolution. In *AI whitepaper*, 2018.
- [33] Indrani Medhi, Aman Sagar, and Kentaro Toyama. Text-free user interfaces for illiterate and semi-literate users. In *2006 international conference on information and communication technologies and development*, pages 72–82. IEEE, 2006.
- [34] Microsoft. Digital agriculture: Farmers in india are using ai to increase crop yields. <https://news.microsoft.com/en-in/features/ai-agriculture-icrisat-upl-india/>, November 2017. (Accessed on 08/30/2020).
- [35] Government of India Ministry of Home Affairs. 2011 census data. <https://www.censusindia.gov.in/2011-Common/CensusData2011.html>. (Accessed on 08/26/2020).
- [36] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 220–229, 2019.
- [37] Shakir Mohamed, Marie-Therese Png, and William Isaac. Decolonial ai: Decolonial theory as sociotechnical foresight in artificial intelligence. *Philosophy & Technology*, pages 1–26, 2020.
- [38] Anahita Mukherji. The cisco case could expose rampant prejudice against dalits in silicon valley. <https://thewire.in/caste/cisco-caste-discrimination-silicon-valley-dalit-prejudice>. (Accessed on 08/14/2020).
- [39] Geoff Mulgan and Vincent Straub. The new ecosystem of trust: how data trusts, collaboratives and coops can help govern data for the maximum public benefit | nesta. <https://www.nesta.org.uk/blog/new-ecosystem-trust/>. (Accessed on 08/21/2020).
- [40] Vasudha Narayanan. Water, wood, and wisdom: Ecological perspectives from the hindu traditions. *Daedalus*, 130(4):179–206, 2001.

- [41] Lawrence A Palinkas, Sarah M Horwitz, Carla A Green, Jennifer P Wisdom, Naihua Duan, and Kimberly Hoagwood. Purposeful sampling for qualitative data collection and analysis in mixed method implementation research. *Administration and policy in mental health and mental health services research*, 42(5): 533–544, 2015.
- [42] C. Rangarajan, S. Mahendra Dev, K. Sundaram, Mahesh Vyas, and K.L Datta. Report of the expert group to review the methodology for measurement of poverty. Technical report, Government of India Planning Commission, 2014.
- [43] World Bank Human Development Unit South Asia Region. People with disabilities in india: From commitments to outcomes.
- [44] Sarah T Roberts. Digital refuse: Canadian garbage, commercial content moderation and the global circulation of social media’s waste. *Wi: journal of mobile media*, 2016.
- [45] RT. ‘whoever leads in ai will rule the world’: Putin to russian children on knowledge day — rt world news. <https://www.rt.com/news/401731-ai-rule-world-putin/>, September 2017. (Accessed on 09/20/2020).
- [46] Anouk Ruhaak. Mozilla foundation - when one affects many: The case for collective consent. <https://foundation.mozilla.org/en/blog/when-one-affects-many-case-collective-consent/>. (Accessed on 08/21/2020).
- [47] Nithya Sambasivan, Garen Checkley, Amna Batool, Nova Ahmed, David Nemer, Laura Sanely Gaytán-Lugo, Tara Matthews, Sunny Consolvo, and Elizabeth Churchill. "privacy is not for me, it’s for those rich women": Performative privacy practices on mobile phones by women in south asia. In *Fourteenth Symposium on Usable Privacy and Security ({SOUPS} 2018)*, pages 127–142, 2018.
- [48] Nithya Sambasivan, Amna Batool, Nova Ahmed, Tara Matthews, Kurt Thomas, Laura Sanely Gaytán-Lugo, David Nemer, Elie Bursztein, Elizabeth Churchill, and Sunny Consolvo. "they don’t leave us alone anywhere we go" gender and digital abuse in south asia. In *proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2019.
- [49] David R Thomas. A general inductive approach for analyzing qualitative evaluation data. *American journal of evaluation*, 27(2):237–246, 2006.
- [50] Kentaro Toyama. *Geek heresy: Rescuing social change from the cult of technology*. PublicAffairs, 2015.
- [51] Tunisia. National ai strategy: Unlocking tunisia’s capabilities potential. In *AI workshop.*, 2018.
- [52] Mazar Ullah. Court told design flaws led to bhopal leak | environment | the guardian. <https://www.theguardian.com/world/2000/jan/12/1>. (Accessed on 08/21/2020).
- [53] Jayapal website. Jayapal joins colleagues in introducing bicameral legislation to ban government use of facial recognition, other biometric technology - congresswoman pramila jayapal. <https://jayapal.house.gov/2020/06/25/jayapal-joins-rep-pressley-and-senators-markey-and-merkley-to-introduce-legislation-to-ban-government-use-of-facial-recognition-other-biometric-technology/>, June 2020. (Accessed on 07/30/2020).
- [54] Virginius Xaxa. Tribes and social exclusion. *CSSSC-UNICEF Social Inclusion Cell, An Occasional Paper*, 2:1–18, 2011.

A Methodology details

We conducted qualitative interviews with 36 expert researchers, activists, and lawyers working closely with marginalised Indian communities at the grassroots. Expert interviews are a qualitative research technique used in exploratory phases, providing practical insider knowledge, surrogacy for a broader community, and importantly [15]. Our respondents were chosen from a wide range of areas to create a holistic analysis of algorithmic power in India. Respondents came from Computer Science (11), Activism (9), Law and Public Policy (6), Science and Technology Studies (5), Development Economics (2), Sociology (2), and Journalism (1). All respondents had 10-30 years of experience working with marginalised communities or on social justice. Specific expertise areas included caste, gender, labour, disability, surveillance, privacy, health, constitutional rights, and financial inclusion. 24 respondents were based in India, 2 in Europe, 1 in Southeast Asia, the rest in the USA; 25 of them self-identified as male, 10 as female, and 1 as non-binary. In conjunction with qualitative interviews, we conducted an analysis of various algorithmic deployments and emerging policies in India, starting from 2009. We identified and analysed various Indian news publications (*e.g.*, TheWire.in, Times of India), policy documents (*e.g.*, NITI Aayog, Srikrishna Bill), grassroots fora (*e.g.*, Dalit Camera, Dalitality), and prior research on automation in India.

We recruited respondents via a combination of reaching out directly and personal contacts, using purposeful sampling [41]—*i.e.*, identifying and selecting experts with relevant experience—iterative until saturation. We conducted all interviews in English (preferred language of participants). Respondents were compensated for the study (giftcards of 100 USD, 85 EUR, and 2000 INR). Interviews lasted an hour each, and were conducted using video conferencing. Transcripts were coded and analyzed for patterns using an inductive approach [49]. From a careful reading of the transcripts, we developed categories and clustered excerpts, conveying key themes from the data. Two team members created a code book based on the themes. The three power structures that we describe below were then developed and applied iteratively to the codes.

We took great care to create a research ethics protocol to protect respondent privacy and safety, especially due to the sensitive nature of our inquiry. During recruitment, participants were informed of the purpose of the study, the question categories, and researcher affiliations. Participants signed informed consent acknowledging their awareness of the study purpose and researcher affiliation prior to the interview. At the beginning of each interview, the moderator additionally obtained verbal consent. We stored all data in a private Google Drive folder, with access limited to our team. To protect participant identity, we deleted all personally identifiable information in research files. We redact identifiable details when quoting participants. Every respondent was given the choice of default anonymity or being included in acknowledgements. All co-authors of this paper work at the intersection of under-served communities and technology, with backgrounds in HCI, critical data studies, and algorithmic fairness; some of us have had grassroots commitments with marginalised Indian communities for over a decade. The first author constructed the research approach, and both the first and second author moderated interviews. All researchers were involved in the research framing, data analysis, and synthesis. Three of us are Indian and two of us are Caucasian. All of us come from privileged positions of class and/or caste.

B Axes of Discrimination in India

In Table 1, we describe some of the prominent axes of injustice in India that were brought up in our interviews, enriched through secondary research and statistics from authoritative sources, to substantiate attributes and proxies to substantiate attributes and proxies.

Sub-groups, Proxies and Harms
<p>Caste (17% Dalits; 8% Adivasi; 40% Other Backward Class (OBC))[35]</p> <ul style="list-style-type: none"> • Societal harms: Human rights atrocities. Poverty. Land, knowledge and language battles [54, 10, 26]. • Proxies: Surname. Skin tone. Occupation. Neighborhood. Language. • Tech harms: Low literacy and phone ownership. Online misrepresentation & exclusion. Accuracy gap of Facial Recognition (FR). Limits of Fitzpatrick scale. Caste-based discrimination in tech ([38]).
<p>Gender (48.5% female)[18]</p> <ul style="list-style-type: none"> • Societal harms: Sexual crimes. Dowry. Violence. Female infanticide. • Proxies: Name. Type of labor. Mobility from home. • Tech harms: Accuracy gap in FR. Lower creditworthiness score. Recommendation algorithms favoring majority male users. Online abuse and 'racey' content issues. Low Internet access.
<p>Religion (80% Hindu, 14% Muslim, 6% Christians, Sikhs, Buddhists, Jains and indigeneous) [35]</p> <ul style="list-style-type: none"> • Societal harms: Discrimination, lynching, vigilantism, and gang-rape against Muslims and others [7]. • Proxies: Name. Neighborhood. Expenses. Work. Language. Clothing. • Tech harms: Online stereotypes and hate speech, <i>e.g.</i>, Islamophobia. Discriminatory inferences due to lifestyle, location, appearance. Targeted Internet disruptions.
<p>Ability (5%–8%+ persons with disabilities) [43]</p> <ul style="list-style-type: none"> • Societal harms: Stigma. Inaccessible education, transport & work. • Proxies: Non-normative facial features, speech patterns, body shape & movements. Use of assistive devices. • Tech harms: Assumed homogeneity in physical, mental presentation. Paternalistic words and images. No accessibility mandate.
<p>Class (30% live below poverty line; 48% on \$2–\$10/day)[42] [35]</p> <ul style="list-style-type: none"> • Societal harms: Poverty. Inadequate food, shelter, health, & housing. • Proxies: Spoken & written language(s). Mother tongue. Literacy. Feature / Smart Phone Ownership. Rural vs. urban. • Tech harms: Linguistic bias towards mainstream languages. Model bias towards middle class users. Limited or lack of internet access.
<p>Gender Identity & Sexual Orientation (No official LGBTQ+ data)</p> <ul style="list-style-type: none"> • Societal harms: Colonial law 377 traces. Discrimination and abuse. • Proxies: Gender declaration. Name. • Tech harms: FR "outing" and accuracy. Gender binary surveillance systems (<i>e.g.</i>, in dormitories). M/F ads targeting. Catfishing and extortion abuse attacks.
<p>Ethnicity (4% NorthEast) [35]</p> <ul style="list-style-type: none"> • Societal harms: Racist slurs, discrimination, and physical attacks. • Proxies: Skin tone. Facial features. Mother tongue. State. Name. • Tech harms: Accuracy gap in FR. Online misrepresentation & exclusion. Inaccurate inferences due to lifestyle, <i>e.g.</i>, migrant labor.

Table 1: Axes of potential ML (un)fairness in India: sub-groups, proxies, and harms.