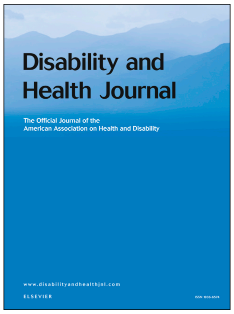

Machine learning to improve analysis of electronic health data on disability and health: an untapped opportunity for health inequities research

Sara Rotenberg, DPhil, Munib Mesinovic, MSc, Eirini-Christina Saloniki, PhD, Shanquan Chen, PhD, Rosalind Raine, PhD, Hannah Kuper, ScD

PII: S1936-6574(25)00246-8

DOI: https://doi.org/10.1016/j.dhjo.2025.102017

Reference: DHJO 102017

To appear in: *Disability and Health Journal*

Received Date: 20 June 2025

Revised Date: 23 October 2025

Accepted Date: 5 December 2025

Please cite this article as: Rotenberg S, Mesinovic M, Saloniki EC, Chen S, Raine R, Kuper H, Machine learning to improve analysis of electronic health data on disability and health: an untapped opportunity for health inequities research, *Disability and Health Journal*, https://doi.org/10.1016/j.dhjo.2025.102017.

Machine learning to improve analysis of electronic health data on disability and health: an untapped opportunity for health inequities research

Sara Rotenberg, DPhil,[1*] Munib Mesinovic, MSc[2], Eirini-Christina Saloniki, PhD[3], Shanquan Chen, PhD[1], Rosalind Raine, PhD[4] , Hannah Kuper, ScD[1]

[1] International Centre for Evidence in Disability, London School of Hygiene and Tropical Medicine, London, UK

[2] Department of Engineering Science, University of Oxford, Oxford, UK

[3] Global Business School for Health, University College London (UCL), London, UK

[4] Department of Primary Care and Population Health, University College London (UCL), London, UK

***corresponding author:** Dr. Sara Rotenberg (sara.rotenberg@lshtm.ac.uk)

Abstract word count: 147

Manuscript word count: 2961

References: 34

Funding Statement

HK and SR are funded by HK's NIHR Global Research Professorship (NIHR301621). SR and SC are funded by the PENDA project from the Foreign, Commonwealth and Development Office. MM is funded by the Rhodes Trust. ECS and RR are funded by the National Institute for Health and Care Research Applied Research Collaboration North Thames. The views expressed in this publication are those of the authors and not necessarily those of the National Institute for Health and Care Research or the Department of Health and Social Care.

Authors' contributions statement

SR designed the outline of the paper and wrote the first draft. MM, ECS, and SC wrote specific sections or provided specific case study examples. All authors contributed to the writing and reviewed and agreed the final version.

Abstract (147 words)

Electronic Health Records (EHRs) are a leading source of epidemiological data, but often lack standardised disability information. This gap hampers our ability to analyse the full scope of health inequities faced by people with disabilities. Current approaches to identify disability within EHRs have limitations because of inadequate proxies for disability or issues linking data sources. Machine learning (ML) offer unprecedented opportunities to create disability markers within EHRs, such as through unsupervised learning to classify disability groups and Natural Language Processing to extract relevant information from clinical notes. These methods have the potential improve disability-disaggregated analyses within EHRs to uncover patterns and provide a more comprehensive understanding of healthcare pathways and outcomes for people with disabilities. Leveraging these approaches to improve disability data in EHRs is a critical step towards improving health inequities research, but emphasise the importance of adhering to ethical guidelines and validating these new approaches.

Background

Globally, there are 1.3 billion people with disabilities, constituting 16% of the population.[1] People with disabilities experience a range of health inequities across the world, including a 13.8-year gap in life expectancy.[2] These gaps arise partially due to inequitable access to healthcare, which manifest as barriers across the healthcare journey for people with disabilities —from a lack of personal health literacy and awareness to unaffordable health services and poor quality care from untrained health workers.[3] Yet, the full scope of these inequities is masked by the sparsity of disability and health data. For example, there are vast gaps in literature around prevalence, prevention, and treatment of common infectious diseases like HIV/AIDS, tuberculosis, and malaria amongst people with disabilities because of a lack of data.[4] Indeed, COVID-19 illuminated the vast inequities in hospitalisation and mortality for people with disabilities, but only because certain data linkage restrictions were lifted to develop a measure of disability in EHRs.[5]

EHRs are heralded as the next frontier for health inequities research, but they rarely include disability questions or markers, making this population unidentifiable within large health data sources.[6] A variety of approaches have been tried to disaggregate EHR data by disability, including linkages to other sources or developing code lists based on certain definitions of disability, but these have limitations and/or are not well-validated.[6, 7] Countries are increasingly enhancing their capacity to collect data on disability in EHRs, such as a new standard coming into effect in the US in 2026 [8] and a disability data asset in Australia [9], but there remain multiple data sources where new data collection or linkages are not possible and so novel methods could be used to close this gap.

This commentary will describe how efforts have been made to improve disability data in EHRs, but the limitations experienced up until now that have hampered our ability to make full use of these data sources.[7] Consequently, we will describe how machine learning (ML) offers opportunities to classify people with disabilities based on patterns from EHRs. ML learns from data to improve performance, such as learning from EHRs to classify groups [10]. We will discuss current challenges with these approaches, focusing on measuring disability; the need for large, linked datasets; and emerging ethical issues. We will make recommendations on further actions to improve disability identification within existing data sources using these novel methods.[10]

Disability definitions and measures

There are different models of disability, and approaches to measure it. However, most countries align with the UN Convention on the Rights of Persons with Disabilities definition and define disability as “long-term physical, mental, intellectual or sensory impairments which in interactions with various barriers may hinder full and effective participation in society on an equal basis with others”.[11]

Disability can be measured several different ways, depending on the definition used. First, it can be self-reported through a simple yes and no question (or set of questions), though this requires individuals to self-identify as disabled. For instance, the 2021 England census asked: "Do you have any physical or mental health conditions or illnesses lasting or expected to last 12 months or more?". If they answered yes, a further question "Do any of your conditions or illnesses reduce your ability to carry out day-to-day activities?" was presented.[12] Second, various questionnaires have been developed to assess self-reported functional difficulties across various functional domains. For example, the Washington Group on Disability Statistics have several question sets, ranging from six to 36 questions to assess self-reported functional impairment across different domains.[13] Other tools to assess disability may include medical assessments (e.g. visual impairment level), or questions on eligibility for disability-based services (e.g., social protection programmes). While all these tools have their strengths, they vary in the conceptual models of disability they measure, their ability to capture the full spectrum of disability, and have trade-offs on specificity, sensitivity, and length of questions.[14, 15]

Importantly, there are no standard questions to assess disability in EHRs, unlike other protected characteristics like ethnicity or gender. Self-reporting of key demographic data, such as disability, is arguably the gold-standard, but this is not presently done in most EHRs. Some health systems (e.g., UK and some US hospital systems) have begun asking about reasonable adjustments often required for people with disabilities such as need for sign language interpreter, height-adjustable exam table, or longer appointment times, [16-19], but this is not a standardised tool or universally collected in all EHRs. Adding to this concern, there are currently no widely-accepted question sets on disability that would be feasible or appropriate for clinical use, as existing tools may be lengthy for clinical settings, or do not cover the full spectrum of disability (e.g., Washington Group Questions).[14, 15] Where disability questions are being asked as part of a medical history, they are rarely updated and unlikely to capture changes in an individual's status over time. Health conditions may be used as a proxy for disability (e.g., cerebral palsy, psychosis), but these also do not map clearly onto the full definition of disability as the diagnosis alone does not indicate the level of functioning and participation restrictions. Thus, disability data are rarely available within EHRs, which would allow tracking of health inequities by disability status.

Current approaches to gather disability data from within electronic health records

There are currently three main approaches to identify disability within EHRs. First, algorithms have been developed to create a marker of disability through an author-defined list of disability-related ICD-9 codes[7]. While these markers provide a good proxy for disability, they are often focussed on particular impairment types (e.g., physical impairment or vision impairment) and only partially validated within a team of researchers and clinicians. [20] Another algorithmic approach linked health insurance claims with self-reported disability data, but insurance claims are not as fulsome sources of data as a patient's full EHR and the validation was poor.[21, 22]

Second, registers have been developed for specific conditions (e.g., vision impairment, intellectual disability) and these can be integrated into EHRs. For example, people with learning disability (intellectual disability) in the UK are eligible to join the learning disability register, which has been used to identify and provide specific services for this population, such as annual health check-ups. The integration into the EHR has also enabled the conduct of analyses by disabilities to understand the health inequities experienced by this population. [23] However, these groupings only work for very specific impairment groups (e.g., visual impairment or intellectual disabilities), and rely on voluntary enrolment (e.g., intellectual disability) or certification by a clinician (e.g., visual impairment), and so coverage is likely to be low.

Third, reasonable adjustments may serve as a proxy for disability. For example, in the UK, health and social care workers can add a 'reasonable adjustments flag' to signal that a person with disabilities may need specific accommodations, such as having longer appointment times for people with intellectual disabilities or height-adjustable examination tables for wheelchair users, and several initiatives in the US request self-reported data [16, 18]. However, these flags and initiatives have relatively low uptake, awareness, and consistency varies significantly by provider, making it a poor marker for disability status in its current form.

An alternative promising approach is to link EHRs with external sources of disability data, such as census data, health insurance, or external registries. For example, in England, individuals' self-reported disability status from the census was linked to EHRs to understand COVID-19 mortality amongst people with disabilities[5]. Similarly, in Wales, EHRs, education and social data were linked to assess whether health checks for adults with intellectual disability contribute to lower mortality for this population group.[24] In South Korea, the national disability register is routinely linked to cancer, and health insurance claims databases to assess health outcomes such as cancer diagnosis, treatment, and outcomes by disability status.[25] However, linkage with other administrative datasets, such as disability benefit claims, often faces technical and legal complexity and these sources may have their own biases because of eligibility and accessibility of assessments narrow those included in the register.

Potential approaches for machine learning

ML can provide an automated alternative approach to produce disability groupings. They can thereby help to optimise current data linkage approaches and offers innovative new ways to link previously unlinked EHR systems. This could be achieved by a machine learning model using an EHR linked to the anchor dataset (e.g., disability register or benefit data) that has disability data and serves as the definition of disability for the model. This two-stage learning set-up is a supervised model that would first learn disability status labels from existing data sources (e.g., census or disability register). The model could learn the ICD-10 or SNOMED codes associated with disability in a less-biased way than author-generated groupings, since it uses an independently validated disability measure. This creates a classifier that learns patterns in diagnostic codes associated with disability. The second stage would be to apply this trained model to other EHRs without disability data or the ability to link to disability data sources to identify individuals likely to have disabilities for the purposes of epidemiological research. There are already several examples of this approach in practice for specific impairment groups. For example, US electronic health record data on older adults were used to identify individuals with functional impairment by linking collected, self-reported functional status to EHRs to predict which clinical variables were associated with functional impairment. While this example provides a strong basis for these methods in older adults and is limited to physical impairment, it could be more broadly applied to identify all types of impairments in EHRs.[26] Data linkage and machine learning are not without their challenges, however. For example, coding of health conditions or disability within EHRs may vary by clinician, and linked datasets used as the disability anchor (e.g., disability benefits) may have their own biases, such as overrepresentation of certain impairment types, age groups, or racial backgrounds. However, this approach would allow us to increase the amount of disability data in EHRs using a tested model developed with self-reported disability status, without the need for new large-scale data collection initiatives.

Another rich source of data on disability could be clinic letters and free-text comments from clinicians documenting relevant clinical history, procedures, or considerations for a patient's care. Physicians may include more detailed information about a person's disability, including how their impairment impacts their participation, rather than a mere clinical diagnosis as proposed in the first method. This information could be extracted and classified using Natural Language Processing (NLP), a subgroup of ML methods, which has already demonstrated significant potential in extracting disability-related information from clinical notes in EHRs to identify people with disabilities Recent studies [27] have shown NLP's effectiveness in identifying functional status, disability, and impairment information from unstructured clinical text, though several papers report that this is still challenging in practice. For instance, Agaronnik et al. [28] developed NLP methods to identify patients with chronic mobility disabilities in EHRs, demonstrating the feasibility of using AI to enhance disability data in health records. However, they encountered challenges in variability due to individual clinician practices, language, and coding, and resorted to manual chart reviews. Other studies have shown that NLP can be leveraged to understand the framing of disability in clinical documentation [27], which is important to understand how various discriminatory biases, such as negative attitudes towards people with disabilities, may appear in clinic letters. This process would require substantial training data and careful validation data to ensure models' generalisability across patient populations, but could offer valuable insights into how providers frame clinic letters and write about various health systems' barriers for patients with disabilities. Better understanding these biases could help create better interventions, such as through more tailored healthcare worker training that addresses particular barriers and biases that come up in clinic letters. While NLP may be a useful tool for disability identification and better understanding biases, further research is needed to overcome the feasibility challenges that have hampered widespread use of this methodology to-date.

Practical and ethical considerations to using machine learning for disability identification

ML often requires an immense amount of data for accuracy, as well as strong definition and rules for what they are trying to measure. ML's transformative potential in disability health is therefore fundamentally constrained by its dependence on high-quality, representative data. Neural networks demand not only substantial dataset volumes but also exceptional data integrity, with minimal missing values or noise contamination. This dual requirement creates a particular challenge in disability health, where successful ML deployment necessitates sophisticated linkage of heterogeneous, longitudinal data sources across fragmented healthcare systems. Even when comprehensive datasets are successfully assembled, ethical AI implementation requires confronting multiple, interconnected sources of bias that threaten both model validity and equitable outcomes. Additionally, another central problem to using ML in EHRs for disability identification is the requirement for strong definitions of disability. Yet, disability is difficult to define and there are different models of disability, making it a challenge to use ML methods to create a singular disability measure.

Moreover, bias permeates ML systems through several critical pathways: documentation inconsistencies across healthcare providers, systematic variations in data quality, and fundamental model inaccuracies. These technical challenges manifest as analytical errors and profoundly skewed outcomes, particularly when training data exhibits disproportionate representation whether through overrepresentation or underrepresentation of specific patient demographics, health conditions, or treatment modalities. The downstream consequence is algorithmic discrimination: unfair performance variations across patient subgroups that directly amplify existing healthcare disparities rather than addressing them. [29]

Mitigating these challenges demands a comprehensive, multi-stage intervention strategy. Pre-processing approaches, including sophisticated resampling techniques and strategic data relabeling, can help balance training datasets while preserving clinical validity[30]. During model training, strategies such as adaptive weight adjustment, emphasizing underrepresented classes while de-emphasizing overrepresented ones, can actively counteract bias.[31] Advanced approaches integrate fairness constraints and regularization directly into loss functions, embedding equity considerations into the optimization process itself.[32] Post-deployment validation through interpretability methods like Shapley values enables ongoing bias detection, revealing whether models make decisions based on clinically relevant features or perpetuate discriminatory patterns embedded in historical data.[33] These comprehensive strategies are not merely technical enhancements but fundamental requirements for responsible ML applications in disability studies.

The development of ML systems capable of identifying disability status also raises profound ethical considerations that extend far beyond technical performance metrics, particularly given disability's status as a protected characteristic in most jurisdictions. The capacity to identify individuals with disabilities with high sensitivity and specificity hinges critically on training data quality, data that carries its own inherent biases and representational limitations, particularly along gender, race, and .[34] This is further complicated in an emerging field, like artificial intelligence, where models of fairness and bias are still being defined. Establishing robust, community-accepted definitions of disability becomes essential for minimizing algorithmic bias, especially as these systems would fundamentally alter the demographic granularity of health records through either data linkage or algorithmic enhancement, often without explicit participant consent for research purposes.

This ethical tension must be balanced against the urgent imperative to address pervasive inequities in health access, quality of care, and outcomes faced by people with disabilities, where data limitations currently impede meaningful progress. Strategic access controls such as restricting linked data and algorithms exclusively to researchers investigating disability inequities rather than extending access to commercial entities like insurance companies can help ensure these powerful tools serve their intended equity-advancing purposes. The phenomenon of disability non-disclosure further complicates these ethical considerations. Stigma within healthcare environments leads many individuals to conceal their disability status, rendering them 'invisible' in datasets unless they specifically seek impairment-related care.[19] Additionally, many people with impairments do not self-identify as disabled, meaning ML systems would effectively assign disability labels without explicit consent. These realities underscore the critical importance of responsible data governance, ensuring these identification tools remain accessible exclusively to equity-focused researchers rather than becoming routine clinical decision-support tools without appropriate safeguards.

Technical validation alone is insufficient for ethical implementation. The disability community must be positioned as active partners in determining the acceptability and appropriateness of these methodological approaches. Meaningful engagement throughout the development process through formal advisory committees, structured consultations, and ongoing collaborative partnerships is not merely advisable but essential for ensuring these ML applications serve rather than exploit disabled populations.

Key takeaways and conclusions

In this paper, we have tried to highlight the ways that advancements in ML could be used to identify people with disabilities in existing data sources, like EHRs, to better understand the health inequities experienced by this population. There are several key takeaways for researchers looking to improve disability data within EHRs. First, ML has already been used in other fields to extract data from clinic letters to improve the richness of EHRs, and this could be done for disability. However, this is not a simple task, and requires a vast amount of data for accuracy, a strong definition of disability to guide the ML on what it needs to predict, and engagement with the disability community for acceptability. Second, gaps in data availability for disability may be solved by linking different data sources, but this remains a problem for data-hungry methods since there are limited data sources where this is possible to do. The adequacy of ML models depends on their inputs, and so careful thought is needed to identify the most optimal data sources to test and validate these models technically, as well as adequate involvement of the disability community for ethical validation. Finally, ML has already shown its limitations and capability for error, and so strict adherence to artificial intelligence ethics principles and careful monitoring are needed in developing these models for disability. While countries catch up with collecting these data in EHRs, these tools may be useful to close the gap, but should be carefully monitored and transparently used to showcase the origins and limitations of the algorithm development.

The health inequities faced by people with disabilities require urgent action. Data has the power to catalyse action on closing these preventable inequities, yet the lack of health data masks these inequities and causes this population to be underserved. The rise of ML has the potential to address this challenge, and it is critical that EHR research improves its ability to capture disability data.

Declarations

Ethics approval and consent to participate – not applicable

Consent for publication – not applicable

Availability of data and materials – not applicable

Competing interests

All authors have no competing interest.

Acknowledgements

None

References:

1. World Health Organization: **Global Report on Health Equity for Persons with Disabilities**. In. Geneva: World Health Organization; 2022.

2. Kuper H, Rotenberg S, Azizatunnisa' L, Banks LM, Smythe T: **The association between disability and mortality: a mixed-methods study**. *The Lancet Public Health* 2024, **9**(5):e306-e315.

3. Kuper H, Azizatunnisa L, Gatta DR, Rotenberg S, Banks LM, Smythe T, Heydt P: **Building disability-inclusive health systems**. *Lancet Public Health* 2024, **9**(5):e316-e325.

4. Smythe T, Rotenberg S, Moyo-Chilufya M, Wilber J, Kuper H: **Health inequities among people with disabilities: an umbrella review and evidence synthesis from 139 meta-analyses**. forthcoming

5. Bosworth ML, Ayoubkhani D, Nafilyan V, Foubert J, Glickman M, Davey C, Kuper H: **Deaths involving COVID-19 by self-reported disability status during the first two waves of the COVID-19 pandemic in England: a retrospective, population-based cohort study**. *The Lancet Public Health* 2021, **0**(0).

6. Kaundinya T: **The Need For Disability Documentation In The Electronic Health Record**. *Health Affairs Forefront* 2022.

7. Brown HK, Carty A, Havercamp SM, Parish S, Lunsky Y: **Identifying reproductive-aged women with physical and sensory disabilities in administrative health data: a systematic review**. *Disability and Health Journal* 2020, **13**(3):100909.

8. Department of Health and Human Services: **21st Century Cures Act: Interoperability, Information Blocking, and the ONC Health IT Certification Program**. In., vol. 85: Federal Register; 2020: 25642-25961.

9. **National Disability Data Asset** [https://www.ndda.gov.au/]

10. Rajpurkar P, Chen E, Banerjee O, Topol EJ: **AI in health and medicine**. *Nat Med* 2022, **28**(1):31-38.

11. United Nations: **UN Convention on the Rights of Persons with Disabilities (CRPD)**. In. New York: United Nations; 2006.

12. **Disability variable: Census 2021** [https://www.ons.gov.uk/census/census2021dictionary/variablesbytopic/healthdisabilityandunpaidcarevariablescensus2021/disability]

13. The Washington Group on Disability Statistics: **WG Short Set on Functioning (WG-SS)**. In.; n.d.

14. Landes SD, Swenor BK, Clark MA, Goddard KS, Hall JP, Hermans A, Ipsen C, Karpman M, Kurth NK, Myers A: **A Research Roadmap Toward Improved Measures Of Disability**. *Health Affairs Forefront* 2024.

15. Mont D, Madans J, Goodman N, Mitra S, Morris Z, Ne'eman A, Wittenburg D: **When Designing Disability Survey Questions, Align Measurement To Purpose: A Response To Landes et al**. *Health Affairs Forefront* 2024.

16. Morris MA, Hamer MK, Eberle K, Jensen KM, Wong AA: **Implementation of Collection of Patients' Disability Status by Centralized Scheduling**. *Jt Comm J Qual Patient Saf* 2021, **47**(10):627-636.

17. **Reasonable Adjustments Flag** [https://digital.nhs.uk/services/reasonable-adjustment-flag]

18. Halkides H, James TG, McKee MM, Meade MA, Moran C, Park S: **Spotlighting Disability in a Major Electronic Health Record: Michigan Medicine's Disability and Accommodations Tab**. *JMIR Form Res* 2022, **6**(12):e38003.

19. Buning GE, James TG, Richards B, McKee MM: **Self-Reported Accommodation Needs for Patients with Disabilities in Primary Care**. *Jt Comm J Qual Patient Saf* 2024, **50**(1):59-65.

20. Khoury AJ, Hall A, Andresen E, Zhang J, Ward R, Jarjoura C: **The association between chronic disease and physical disability among female Medicaid beneficiaries 18–64 years of age**. *Disability and Health Journal* 2013, **6**(2):141-148.

21. Ben‐Shalom Y, Stapleton DC: **Predicting disability among community‐dwelling Medicare beneficiaries using claims‐based indicators**. *Health Services Research* 2016, **51**(1):262-281.

22. Palsbo SE, Sutton CD, Mastal MF, Johnson S, Cohen A: **Identifying and classifying people with disabilities using claims data: further development of the Access Risk Classification System (ARCS) algorithm**. *Disability and Health Journal* 2008, **1**(4):215-223.

23. Williamson EJ, McDonald HI, Bhaskaran K, Walker AJ, Bacon S, Davy S, Schultze A, Tomlinson L, Bates C, Ramsay M *et al*: **Risks of covid-19 hospital admission and death for people with learning disability: population based cohort study using the OpenSAFELY platform**. *BMJ* 2021, **374**:n1592.

24. Kennedy N, Kennedy J, Kerr M, Dredge S, Brophy S: **Health checks for adults with intellectual disability and association with survival rates: a linked electronic records matched cohort study in Wales, UK**. *BMJ Open* 2022, **12**(4):e049441.

25. Ko KD, Lee KY, Cho B, Park MS, Son KY, Ha JH, Park SM: **Disparities in health-risk behaviors, preventive health care utilizations, and chronic health conditions for people with disabilities: the Korean National Health and Nutrition Examination Survey**. *Archives of Physical Medicine and Rehabilitation* 2011, **92**(8):1230-1237.

26. Pavon JM, Previll L, Woo M, Henao R, Solomon M, Rogers U, Olson A, Fischer J, Leo C, Fillenbaum G *et al*: **Machine learning functional impairment classification with electronic health record data**. *J Am Geriatr Soc* 2023, **71**(9):2822-2833.

27. Newman-Griffis DR, Desmet B, Zirikly A, Tamang S, Chang CH: **Editorial: Artificial intelligence for human function and disability**. *Front Digit Health* 2023, **5**:1282287.

28. Agaronnik ND, Lindvall C, El-Jawahri A, He W, Iezzoni LI: **Challenges of Developing a Natural Language Processing Method With Electronic Health Records to Identify Persons With Chronic Mobility Disability**. *Archives of Physical Medicine and Rehabilitation* 2020, **101**(10):1739-1746.

29. Chen F, Wang L, Hong J, Jiang J, Zhou L: **Unmasking bias in artificial intelligence: a systematic review of bias detection and mitigation strategies in electronic health record-based models**. *Journal of the American Medical Informatics Association* 2024, **31**(5):1172-1183.

30. Li F, Wu P, Ong HH, Peterson JF, Wei W-Q, Zhao J: **Evaluating and mitigating bias in machine learning models for cardiovascular disease prediction**. *Journal of Biomedical Informatics* 2023, **138**:104294.

31. Lahoti P, Beutel A, Chen J, Lee K, Prost F, Thain N, Wang X, Chi E: **Fairness without demographics through adversarially reweighted learning**. *Advances in neural information processing systems* 2020, **33**:728-740.

32. Berk R, Heidari H, Jabbari S, Joseph M, Kearns M, Morgenstern J, Neel S, Roth A: **A convex framework for fair regression**. *arXiv preprint arXiv:170602409* 2017.

33. Watson D: **Rational Shapley Values**. In: *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency.* Seoul, Republic of Korea: Association for Computing Machinery; 2022: 1083–1094.

34. Newman-Griffis D, Rauchberg JS, Alharbi R, Hickman L, Hochheiser H: **Definition drives design: Disability models and mechanisms of bias in AI technologies**. *arXiv preprint arXiv:220608287* 2022.