

Data, Power and Bias in Artificial Intelligence

Susan Leavy^{1,2}, Barry O’Sullivan^{1,3} and Eugenia Siapera²

¹Insight Centre for Data Analytics

²School of Information and Communication Studies, University College Dublin, Ireland

³School of Computer Science and Information Technology, University College Cork, Ireland

susan.leavy@ucd.ie | b.osullivan@cs.ucc.ie | eugenia.siapera@ucd.ie

Abstract

Artificial Intelligence has the potential to exacerbate societal bias and set back decades of advances in equal rights and civil liberty. Data used to train machine learning algorithms may capture social injustices, inequality or discriminatory attitudes that may be learned and perpetuated in society. Attempts to address this issue are rapidly emerging from different perspectives involving technical solutions, social justice and data governance measures. While each of these approaches are essential to the development of a comprehensive solution, often discourse associated with each seems disparate. This paper reviews ongoing work to ensure data justice, fairness and bias mitigation in AI systems from different domains exploring the inter-related dynamics of each and examining whether the inevitability of bias in AI training data may in fact be used for social good. We highlight the complexity associated with defining policies for dealing with bias. We also consider technical challenges in addressing issues of societal bias.

1 Introduction

The detrimental effects of unfairness and bias in artificial intelligence (AI) on equal rights in society have been well documented. Predictive policing systems trained on historical police records are being increasingly used to forecast criminal behaviour even though the accuracy of this data has been called into question [Richardson *et al.*, 2019]. Racial bias was uncovered in an algorithm used to support healthcare decisions [Obermeyer *et al.*, 2019]. While bio-metric technology such as facial recognition is being increasingly integrated into core security and border control infrastructures, their accuracy is lowest on darker-skinned females [Buolamwini and Gebru, 2018]. In search engines and recommender systems, race and gender discrimination was uncovered across multiple online platforms [Datta *et al.*, 2015; Lambrecht and Tucker, 2019; Noble, 2018]. In each of these cases the effects of bias in AI was to discriminate against those already marginalised in society.

Bias and discrimination in AI follows a long history of injustice resulting from the way people are counted, represented

and classified in data [Perez, 2019; Hickman, 1997]. To those affected by injustice, whether unfair decisions are made by human or machine is not the most pressing issue. It is perhaps for this reason that the focus of analysis in terms of social justice, critical race studies and feminism is not solely restricted to artificial intelligence but algorithms, socio-technical systems and automated decision making.

In parallel, within the field of AI, approaches to ensuring fairness and justice often assume that entirely new theoretical and ethical frameworks are required. However, as is clear in relation to the use of historical police records to predict future crime [Lum and Isaac, 2016], the same critical perspectives and activism that brought social justice in the past are fundamental to preventing bias and discrimination in AI. This paper explores this issue examining how research grounded in social justice, critical race studies and feminism could form the foundation of a new approach to data collection and curation for AI. We also consider technical challenges in addressing issues of societal bias.

2 The Myth of Objectivity

Approaches to the prevention of bias and discrimination in AI have included developing fairness-aware machine learning algorithms, addressing bias in data and modifying learned models. The central aim of this work is to assume the status of objectivity and neutrality. However, there is increasingly widespread acknowledgement of the impossibility of objectivity in data-driven AI [Meredith, 2018; O’Neil, 2016] and of bias as “*an unavoidable characteristic of data collected from human processes*” [Dignum, 2019]. This sentiment was aptly captured by Bartoletti [2020] in her assertion that “*it is time we acknowledge that data is simply not neutral, and, as such, every single decision and action around data is a political one.*” This acknowledgement of the inevitability of bias necessitates a reform of the process of collating and curating training data for machine learning, incorporating perspectives across multiple disciplines.

Bringing to light the extent and complexity of damaging constructs of race and gender and how it is embedded in both contemporary and historical data makes it clear how unavoidable it is that data reflects some form of bias or captures the effects of social injustice. Noble [2018] uncovered a series of cases where data-driven racial profiling of individuals resulted in the repetition of historical injustices against people

of colour through what she termed ‘*technologically redlining*’. How ‘*runaway feedback loops*’ are generated by machine learning algorithms when they are trained on data large datasets that capture the results of a history social injustice and how this exacerbates existing discriminatory practices in society is detailed by Gebru [2019].

The process of data collection and curation for machine learning algorithms, from the standpoint of feminism or critical race studies, can never be objective. Every dataset that aims to represent humans and their behaviour does so in accordance with a world-view or ideology that may be assimilated into an AI system. Data collection and curation for machine learning algorithms is therefore transformed into a political act of curating a world-view that one intends to perpetuate through an AI system.

3 Beyond Bias

In contrast to the increasingly widespread acknowledgement of the impossibility of objectivity in data, most current technical approaches to bias in training data for machine learning aim to re-balance data and render it neutral in relation to protected classes such as race or gender. These include methods for data augmentation, re-sampling, re-weighting, swapping labels and removing dependencies between characteristics [Kamiran and Calders, 2012; Feldman *et al.*, 2015].

In relation to language based datasets, methods to uncover and address bias include the use of word embedding models trained on text to uncover social and intersectional bias [Tan and Celis, 2019]. Stereotypical associations in text were amended by disassociating relationships between entities in the trained models [Bolukbasi *et al.*, 2016] and by swapping the gender of entities in text [Zhao *et al.*, 2018; Park *et al.*, 2018]. However, the evaluation of bias in these studies focus on specific aspects such as stereotypical associations of race and gender in relation to employment or use implicit association tests as evaluation frameworks for bias in the study by Caliskan [2017].

Within narrowly defined contexts, approaches to modifying data have been shown to reduce bias in algorithms. However, grounding the work in critical studies of race, feminist theories and social justice may serve to enrich the frameworks for evaluating fairness in how groups are represented in data. Furthermore, incorporating such perspectives diverts efforts away from a search for bias towards a more comprehensive interrogation of the power structures and ideologies of gender, race and social class that may be captured in the data.

In assessing the ideology underlying a dataset, the critical framework proposed by D’Ignazio and Klein [2020] presents a comprehensive approach to the evaluation of how gender is represented in data, grounded in intersectional feminist thought. The authors develop feminist principles for data collection asserting that “*data are not neutral or objective.*” This approach if embedded into the data collection process for machine learning would highlight potentially problematic patterns in existing data that may lead to discriminatory AI systems. It may also work towards the transformation of the process to one involving the curation of data collections that intentionally reflect an intersectional feminist representation

of gender.

By bridging critical theories of race, social justice or feminist theory with training machine learning algorithms in this way, less focus may be placed on attempting to ‘fix’ specific issues of bias in datasets that contain embedded problematic ideologies. Rather, a dataset may be designed and selected in order to capture a representation of the kind of concepts of race, gender and social class that should be perpetuated in an AI algorithm.

4 Power in Data

One response to bias in data resulting from an under-representation of particular groups, is to collect more data from that group of people. This has led to questions of unethical data gathering practices with the goal of balancing datasets [Hawkins, 2018]. Benjamin [2019] examines the complex socio-political ramifications of this in the context of race, questioning “*what does it mean to be included, and hence more accurately identifiable, in an unjust set of social relations?*.” This historical legacy of inequitable treatment of the poorer in society in relation to data collection is described by Eubanks [2018] as what she termed the ‘*Digital Poorhouse.*’ In attempting to balance data sets for training machine learning algorithms, it is crucial then to consider the social and political effects of data collection for different groups in society.

The central importance of fair data in preventing discrimination in AI, highlights the central role played by curators of data for machine learning algorithms. The cognitive labour of classifying and labelling AI training data is increasingly conducted by underpaid women of colour from the global south [Crawford and Joler, 2018; D’Ignazio and Klein, 2020]. However, with the increased recognition that bias and discrimination in AI emanates largely from data, along with the deconstruction of the myth of algorithmic objectivity, it should follow that the human labour and value involved in data curation for AI will become more visible. It also emphasises the case for what Benjamin [2019] describes as the ‘*democratization of data*’ through the inclusion of all perspectives, especially from those most vulnerable to discrimination, into the design of technology.

5 Governing Artificial Intelligence

The protection of autonomy and the right to self-determination of every individual is central to research and activism at the intersection of technology and social justice. By virtue of its capacity to learn, AI raises foundational ethical questions in relation to human autonomy and introduces new dimensions in terms of accountability for discrimination. Ensuring humans are accountable for decisions made by AI systems is key to ethical AI [O’Neil, 2016]. To reflect the foundational nature of such questions, governing bodies are developing new frameworks for ethical, trustworthy and responsible AI primarily grounded in a human rights perspective [EU-HLEG-AI, 2019; OECD, 2019; US-Govt, 2019]. However, given the accelerating pace of the influence of AI in society there is also an urgency on the implementation of measures to prevent discrimination.

In addressing the urgent need for governance, organisations are examining the use of existing data governance policies in the context of the collection and curation of machine learning training data [Adamson *et al.*, 2019; ISO, 2017]. There are also calls for standards on the kind of inferences that AI algorithms can make based on an individual’s data [Wachter and Mittelstadt, 2019]. However, the implementability of data governance regulations to ensure bias-free data remains a challenge. The central aim of ensuring bias-free data contrasts with the increasingly accepted position that data, particularly in relation to human behaviour, is inevitably biased. Additionally, within AI literature there is a range of different definitions of fairness and methods of representing the concept mathematically [Mehrabi *et al.*, 2019]. Departing from the concept of bias, which implies the possibility of bias-free data and grounding data regulation instead in concepts of discrimination derived from well established legal definitions, may contribute towards unifying research across the disciplines of AI, policy and social justice.

6 Technical Challenges and Bias

A considerable challenge in dealing with bias in AI generally, and in machine learning in particular, is that there can be different definition of bias or what it means to be fair. We consider here the often discussed COMPAS system for predicting recidivism.¹ Tables 1 and 2 are based on an analysis presented by Sumpter [2018] that considered whether or not COMPAS exhibits racial bias. The answer is that it depends on what one defines as bias.

Kleinberg *et al.* [2017] considered three different fairness conditions that capture what is means for an algorithm to be fair and avoid exhibiting bias. The first of these is that the algorithm should be *well calibrated*. Essentially, this means that the proportion of people that are positive in a population should be equal to the proportion of people that are positive in each subgroup of the population. Comparing Tables 1 and 2 we see the that $1369/2174 = 63\%$ of African American defendants classified as high risk, while $505/854 = 59.1\%$ of White defendants were classified as such. This data is well-calibrated, suggesting the system does not exhibit a racial bias, since the proportion of high-risk people in each population is predicted to be roughly equal.

The second (and third) condition relates to balancing for the positive (resp. negative) class. If this condition were to be violated it would mean that the likelihood that a positive (resp. negative) instance in one population is more likely to be identified than in the other. For example, as Sumpter argues, Tables 1 and 2 show that $2174/3615 = 60\%$ of African Americans in COMPAS are considered higher risk, while this only $854/2464 = 35\%$ for Whites; note that $1901/3615 = 53\%$ of African Americans reoffended while this was $966/2454 = 39\%$ for White, suggesting that African Americans actually reoffended less than predicted while the opposite is true for Whites. This suggests a strong racial bias.

Considering the mistakes that COMPAS makes for each racial group, Sumpter goes on to show that $805/1714 = 47\%$

¹<https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>

Table 1: Recidivism rates in the COMPAS system for African American defendants (via [Sumpter, 2018]).

African American	High Risk	Low Risk	Total
Reoffended	1,369	532	1,901
Didn’t reoffend	805	990	1,714
Total	2,174	1,522	3,615

Table 2: Recidivism rates in the COMPAS system for White defendants (via [Sumpter, 2018]).

White	High Risk	Low Risk	Total
Reoffended	505	461	966
Didn’t reoffend	349	1,139	1,488
Total	854	1,600	2,454

of African Americans were wrongly predicted to reoffend, as compared with only $349/1488 = 24\%$ of Whites. On the other hand, only $532/1901 = 28\%$ of African Americans who reoffended were wrongly predicted as being lower risk, as compared with $461/966 = 48\%$ of Whites. Again, this suggests a strong racial bias.

In other words, while COMPAS satisfies one of Kleinberg *et al.*’s definition of fairness (calibration), it fails the second and third. This is not unusual. In fact, Kleinberg *et al.* rigorously prove that these three fairness conditions are incompatible except in very rare situations. This implies that one is usually faced with a choice of which bias must be traded-off against another. Eliminating bias from AI systems is often provably impossible.

The challenges in dealing with bias are complex. To address these issues, in addition to the many point made earlier, we must develop a set of AI methods that reason about bias as systems are being developed. It might be helpful to think of bias-aware learning as an optimisation problem in which the properties of the learned concept and the acceptable tradeoff between fairness criteria could be imposed as constraints on the learning process [Bessiere *et al.*, 2009].

7 Conclusion

In considering the challenges of bias in AI, it is important to build upon existing principles and frameworks, e.g. critical intersectional feminist and critical race theories. Dealing with bias in an objective way is very challenging and the possibility of unbiased data is not a sufficient condition for unbiased AI. In reasoning about bias and fairness it is critically important to involve those who might be directly impacted by AI system in the process of designing, curating, testing these systems, and especially in the ‘deconstruction’ of language-based data. It is important that we interrogate notions of accuracy that may lead to more surveillance and data gathering on marginalised groups, and the role of such marginalised groups in data curation is important to acknowledge.

Acknowledgement. This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant number 12/RC/2289-P2 which is co-funded under the European Regional Development Fund.

References

- [Adamson *et al.*, 2019] G. Adamson, J. C. Havens, and R. Chatila. Designing a value-driven future for ethical autonomous and intelligent systems. *Proceedings of the IEEE*, 107(3):518–525, 2019.
- [Bartoletti, 2020] Ivana Bartoletti. *An Artificial Revolution: On Power, Politics and AI*. The Indigo Press, 2020.
- [Benjamin, 2019] Ruha Benjamin. *Race after technology: Abolitionist tools for the new jim code*. John Wiley & Sons, 2019.
- [Bessiere *et al.*, 2009] Christian Bessiere, Emmanuel Hebrard, and Barry O’Sullivan. Minimising decision tree size as combinatorial optimisation. In Ian P. Gent, editor, *Principles and Practice of Constraint Programming - CP 2009, 15th International Conference, CP 2009, Lisbon, Portugal, September 20-24, 2009, Proceedings*, volume 5732 of *Lecture Notes in Computer Science*, pages 173–187. Springer, 2009.
- [Bolukbasi *et al.*, 2016] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems*, pages 4349–4357, 2016.
- [Buolamwini and Gebru, 2018] 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.
- [Caliskan *et al.*, 2017] Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, 2017.
- [Crawford and Joler, 2018] Kate Crawford and Vladan Joler. Anatomy of an ai system. Retrieved from: <https://anatomyof.ai/>, 2018.
- [Datta *et al.*, 2015] Amit Datta, Michael Carl Tschantz, and Anupam Datta. Automated experiments on ad privacy settings. *Proceedings on privacy enhancing technologies*, 2015(1):92–112, 2015.
- [D’Ignazio and Klein, 2020] Catherine D’Ignazio and Lauren F Klein. *Data feminism*. MIT Press, 2020.
- [Dignum, 2019] Virginia Dignum. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Springer International Publishing, 2019.
- [EU-HLEG-AI, 2019] EU-HLEG-AI. High-level expert group on artificial intelligence: Ethics guidelines for trustworthy ai. *European Commission*, 09.04, 2019.
- [Eubanks, 2018] Virginia Eubanks. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press, 2018.
- [Feldman *et al.*, 2015] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and removing disparate impact. In *proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pages 259–268, 2015.
- [Gebru, 2019] Timnit Gebru. Oxford handbook on ai ethics book chapter on race and gender. *arXiv preprint arXiv:1908.06165*, 2019.
- [Hawkins, 2018] Amy Hawkins. Beijing’s big brother tech needs african faces. *Foreign Policy*, 24, 2018.
- [Hickman, 1997] Christine B Hickman. The devil and the one drop rule: Racial categories, african americans, and the us census. *Michigan Law Review*, 95(5):1161–1265, 1997.
- [ISO, 2017] ISO. Iso 26000 social responsibility, 2017.
- [Kamiran and Calders, 2012] Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1):1–33, 2012.
- [Kleinberg *et al.*, 2017] Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPICs*, pages 43:1–43:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [Lambrecht and Tucker, 2019] Anja Lambrecht and Catherine Tucker. Algorithmic bias? an empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science*, 2019.
- [Lum and Isaac, 2016] Kristian Lum and William Isaac. To predict and serve? *Significance*, 13(5):14–19, 2016.
- [Mehrabi *et al.*, 2019] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*, 2019.
- [Meredith, 2018] B Meredith. Artificial unintelligence: How computers misunderstand the world, 2018.
- [Noble, 2018] Safiya Umoja Noble. *Algorithms of oppression: How search engines reinforce racism*. nyu Press, 2018.
- [Obermeyer *et al.*, 2019] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019.
- [OECD, 2019] OECD. Oecd council recommendation on artificial intelligence, 2019.
- [O’Neil, 2016] Cathy O’Neil. *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books, 2016.
- [Park *et al.*, 2018] Ji Ho Park, Jamin Shin, and Pascale Fung. Reducing gender bias in abusive language detection. *arXiv preprint arXiv:1808.07231*, 2018.
- [Perez, 2019] Caroline Criado Perez. *Invisible Women: Exposing data bias in a world designed for men*. Random House, 2019.

- [Richardson *et al.*, 2019] Rashida Richardson, Jason Schultz, and Kate Crawford. Dirty data, bad predictions: How civil rights violations impact police data, predictive policing systems, and justice. *New York University Law Review Online*, Forthcoming, 2019.
- [Sumpter, 2018] David Sumpter. *Outnumbered: From Facebook and Google to Fake News and Filter-bubbles – The Algorithms That Control Our Lives*. 2018.
- [Tan and Celis, 2019] Yi Chern Tan and L Elisa Celis. Assessing social and intersectional biases in contextualized word representations. In *Advances in Neural Information Processing Systems*, pages 13209–13220, 2019.
- [US-Govt, 2019] US-Govt. Guidance for regulation of artificial intelligence applications, 2019.
- [Wachter and Mittelstadt, 2019] Sandra Wachter and Brent Mittelstadt. A right to reasonable inferences: re-thinking data protection law in the age of big data and ai. *Colum. Bus. L. Rev.*, page 494, 2019.
- [Zhao *et al.*, 2018] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Gender bias in coreference resolution: Evaluation and debiasing methods. *arXiv preprint arXiv:1804.06876*, 2018.