

The Authority of "Fair" in Machine Learning

Anonymous Author(s)

ABSTRACT

In this paper, we argue for the adoption of a normative definition of fairness within the machine learning community. After characterizing this definition, we review the current literature of Fair ML in light of its implications. We end by suggesting ways to incorporate a broader community and generate further debate around how to decide what is fair in ML.

1 INTRODUCTION

The recent boom of machine learning (ML) applications has just as quickly given rise to a slew of critics pointing out the harmful capabilities of these systems. In particular, concerns of bias and discrimination are being debated as ML systems for natural language processing [7], judicial sentencing [21], target advertising [13, 34], image classification [1], and facial recognition [9, 35] have all proven their ability to inherit bias and create disparate treatment across groups. Responding to these findings is a body of work that attempts to import considerations of "fairness" into our ML approaches [3, 10, 18, 20].

In this paper, we argue for expanding and deepening our approach to "fairness" in ML practice. Drawing from philosophy and ethics, we offer up a normative account of fairness where "fair" is a property that is both communally derived and context dependent. Using this definition, we highlight three categorical framings through which one may inquire about the "fairness" of an ML system: fairness of a system, fairness of an approach, and fairness of a result. We justify these different framings and then move on to overview contemporary approaches to fairness in the ML literature. We argue the position that the literature has thus far focused on the problem of disparate treatment without much attention to other framings. Making this salient allows us to consider the importance of critically examining "whose version of fair" we privilege in ML moving forward and argue that taking a stance on fairness must be understood as invoking an authority that could be more or less legitimate. We conclude by offering possible pathways forward for the community to broaden its approach to fair ML.

2 THE CONSTRUCTION OF "FAIRNESS"

There was once a time when fairness or "the good" was believed to be a static and essential property that could be derived from divine [27] or rational principle [14]. Determining what was right or fair was a privilege vested in certain authorities and the results were absolute, allowing for no disagreement. This traditional view has changed as modern philosophy and ethics has revised our typological ways of thinking into a normative framing where concepts such as "fair" or "just" are no longer static; rather, they are developed relative to particular communities (who) and contexts (when/how). This is why for modern philosophers such as Richard Rorty [29] or Chantal Mouffe [25], disagreement is part and parcel of a democratic society where we must navigate these tensions in hopes of finding places of commonality from which to move forward.

We are now in an age of *machine action* where algorithms can benefit some individuals but may do so at the cost of harming others. Thus, we must not take the responsibility of implementing fair ML systems lightly. Transitioning fragile and contentious matters of human judgment to trained models must be done with care and forethought. As things stand, any engineer with a data set may codify a notion of fairness into an ML system without allowing for any disagreement or community consensus. In order to avoid slipping back into an essentialist morality where a small elite group decides what makes a machine action fair with no recourse, it is critical we expand the available framings and considerations of "fairness" in ML.

Taking the stance that we should apply our modern ideas of fairness to ML, we offer the following proposition:

Proposition (P1): A machine learning system can only be fair with a contextual justification for the choice of a fairness construct and offering a channel for affected parties to actively assent or dissent to the fairness of the system.

In P1, a "fairness construct" is a definition for what fair is taken to mean in the problem space and the approach used to codify and measure that definition in training and application. What P1 implies is that a team implementing an ML system should have a sense of the viewpoints around fairness in a domain and be prepared to take a stance when choosing an approach. Further, it dictates some mode of disagreement be available. For example, either a) creating a method of algorithmic due process [12] where the fairness of a result can be scrutinized or b) working with an active community to address and resolve disputes over time.

A real-world example where we may require such a definition of fairness comes around building classifiers for predicting mental health issues using social media data. In research, there is a tacit acceptance that using social media data to predict mental health is fair [24, 31]. Models and outputs for this have been peer-reviewed and the system itself appears to pass as fair, or at least acceptable. Adopting the normative standard set by this work, one may believe it is fair for anyone to build a similar system to predict mental health, no matter what. And this is exactly what Facebook did—it was classifying teenagers by their psychological vulnerabilities such as feeling "insecure," "worthless," or "needing a confidence boost" [32]. However, Facebook's practice caused a lot of backlash from its users. The reaction to what such a system looks like in practice should raise a red flag that our questions about fairness must go beyond "was the technical approach fair" or "are the results fair." It is our belief that the community of ML researchers and practitioners should also be asking questions such as "in what context?", "with what dataset?", and "with what objective?" should a system be trained to classify mental health. It is here we might consider that the discussion of Fair ML cannot be constrained to whether the classifier responds equally to similar inputs. It is our position that ML practitioners must not skirt responsibility in questioning the ethics of what they build so long as some minimum equality criterion is met.

2.1 Contextualizing Fairness

In order to clarify our position on the role fairness should play in ML, we offer up three categorical questions one might use to frame an inquiry around the fairness of an arbitrary ML system: **(Q1)** "Is it fair to make X ML system?"; **(Q2)** "I want to make X ML system, is there a fair technical approach?"; and **(Q3)** "I made X ML system, are the results fair?". Each question requires a different set of considerations to arrive at an answer.

Q1 is asking whether a particular problem, in general, is fit to be approximated or decided upon by an ML system. Using P1, this kind of question requires us to consider the sentiments of the communities it would effect and whether we have a sense of what a fair automation of the task would look like. An immediate response may be, "well anything is worth a try if there's a reasonable data set and approach," but we will come back to why that is a specious assumption.

Q2 relies on a series of methodological questions perhaps best suited for experts. The fairness of an approach might rely on knowing what a good sample space might look like and the potential biases in historical samples. Answering this question would require an inquiry around whether the set of available features is a close and fair approximation of what we want to predict. Finally, we would hope a practitioner

can weigh in on different trade-offs needing consideration for the choice of a training regime (e.g., [22, 43]).

Q3 gets at the question of treatment. That is, do the outputs of the algorithm actually correspond to what would constitute a fair response by a human. Answering Q3 requires us to run tests against the model and unpack the algorithm to determine qualities such as what input features were considered important, whether or not different groups have equal chance for misclassification, and whether any variables were acting as proxies for protected variables.

In light of these multiple framings, we take the position that no single framing nor problem space should dominate the realm of what we may consider fair ML practice. While we recognize there will always be organizational ethics that must be further considered beyond what an ML practitioner can influence, we side with an interpretation of engineering ethics similar to that espoused by Langdon Winner [39] that the technology artifact itself has politics and is thus appropriate for normative evaluation. That is, he rejects that all ethics around technological artifacts are socially determined and argues that history has shown us that the artifacts themselves carry political and ethical weight. We now offer a few reasons why the advent of machine learning may embolden such a stance.

2.2 Why "Fair" Matters

In line with Winner's position, we should not see an ML model as a blank slate that can only be evaluated after appropriation by some organization. Not only might an ML artifact mediate ethically charged situations, but further it is an artifact carrying some amount of agency. From this vantage, "fair" ML is a recasting of the very idea of fair action in the human sense. Though it may be perceived that adding a "human in the loop" could solve our issue of ethical machine autonomy, research shows that "users may be prone to place an inordinate amount of trust in black box algorithms that are framed as intelligent" [33]. Meaning that even when an ML model is not acting autonomously, it is causing normative sway that is not neutral.

Further, we must consider that as actions coordinated by an ML model intervene in more of our lives, these actions are not always welcomed or requested [4, 5, 36]. As Frank Pasquale points out in "The Black Box Society," how we are categorized through data affects how we will be treated [26]. Grounding this thought, we ask whether someone who has never requested therapy, counseling, or any sort of diagnosis should be considered "open game" for a mental health classification. There is not an easy answer to this which begs the original question of whether any ML model that classifies mental health on social media data is fair.

Finally, we want to point to the fact that often the choice of a training objective is contentious in and of itself. There

are some cases where the objective has a clear consensus, say in the case of classifying radiology images by whether a cancerous tumor is observed. There is little dispute that the goal here is to be as accurate as possible at identifying cancer. On the other hand, determining whether an individual convicted of a crime might be a repeat offender is likely to solicit a lot of disagreement. Re-appropriating terminology from Richard Rorty [30], we distinguish between these two cases as *normal* and *abnormal* objectives, respectively. Meaning the objective of some ML tasks have a very clear grounds for consensus (normal) and others are highly disputable (abnormal). Due to the fact that ML could be applied to either kind of task, we believe this consideration elevates the level to which an ML model may be considered political or up for normative evaluation.

3 CURRENT STATE OF FAIR ML

In light of our argument that questions of fairness operate contextually and that the advent of ML elevates our need for normative evaluation of technology artifacts, we move on to apply our contextual framings of fair to the contemporary Fair ML literature. We believe progress has been made, but mostly within the scope of framings Q2 and Q3 and almost exclusively employing some variant of "preventing disparate impact" as the definition of fairness.

3.1 Is it Fair to Make X ML system?

Outside of broad critical reflections on the use of technical systems [8, 16], there is yet to be much work characterizing grounds for why an ML system may or may not be a fair approach for a particular problem space. One might anticipate, given a normative definition of fairness, that questions of whether any data could ever approximate certain problems would provide grounds for healthy cross-disciplinary debate. However, "If we can, we ought to," is often treated as an unstated premise for technological development. An exceptional case, is the contemporary debate around autonomous weapons where thousands of scientists have supported a total suppression of development in this area [17].

In the Fair ML literature, a recent example of researchers asking more general questions about fairness is [6]'s evaluation of the ethical implications of ML for autonomous experimentation. Recognizing that practitioners have largely ignored established human-subject research guidelines laid out by the Common Rule and Belmont Principles, the authors argue that automated experimentation may cause harms from privacy-violating inferences and exposing users to less-than-ideal outcomes by being part of the experiment. While falling short of actually questioning whether automated experimentation should be allowed at all, they suggest the adoption of an external review process in light of the fact that it's intractable to obtain consent from each user due

to the complexity of the systems being tested. One might construe their high-level question as "Is it fair to make automated experimenting ML systems?" and their answer as "Maybe, but we should have external oversight."

Though we would be interested in a broader debate about the fairness of research ethics using ML systems, the conclusion drawn supports a more normative evaluation of fairness in this realm. Specifically, the idea that external review may be needed hinges on ideas of authority and context; namely, should we give engineers blanket authority to experiment on users? Akin to P1 above, research ethics are such that a review board should check your experimental standards and participants must be given certain rights. Thus, we agree with the suggestion of external review and further that certain autonomous experiments likely should not be done. Through our lens we would argue many other areas (e.g., ML for biometric inference, ML for emotional persuasion) are ripe for debate around the limitations of what systems are fair to implement.

3.2 I want to make X ML system, is there a fair technical approach?

Surveying the ML literature, we see three categorical trends signifying answers this question. The first grouping is research related to interpretability or "white box models" [37, 38, 40]. That is, this body of research approaches fairness by developing training methods that aim to produce interpretable ML models. Whether the model is fair becomes a matter of whether an explanation of the results is interpreted as fair. This intersects with the EU's upcoming adoption of "the right to an explanation" law and researchers' calls for algorithmic due diligence [12]. In our view, this approach has the most affinity with our P1 fairness definition given that it opens up the possibility for models to be interpreted by various subjects (allowing for contextual considerations) and sets forth future possibilities for recourse and disagreement. The limitations of this approach are 1) this approach is not yet feasible for more complex algorithms (ie, most current trends in the ML field) and 2) fairness is bound by the ability for a subject to meaningfully understand and act on the interpretation as the interpretability guarantees nothing about the fairness of the algorithm itself.

Our second categorization includes attempts to resolve disparate treatment concerns by developing statistical independence between predictions and protected categories. These approaches include methods to satisfy fairness by enforcing robust sampling across groups [10], minimizing the difference in misclassification rates across groups [42], and training models where protected variables are neither explicitly nor latently used [23]. All of these methods define fairness in relation to treatment across known protected

classes and give authority to an engineer (or perhaps partnered legal advisor) to *a priori* determine which variables are protected. This is a limitation due to the fact that some problem spaces may not have obvious or measurable categories that deserve protection (e.g., should we target someone for prescription drugs at a inferred moment of vulnerability?). A further limitation to this class of approaches shown by ML researchers is that there are inherent trade-offs between a) well-calibrated models, b) parity between groups in the positive class, and c) parity between groups in the negative class [22].

A third category of work involves decision-making algorithms that attempt to construct fair metrics for optimal choices. These attempts start with the adoption of a quality function that evaluates decisions and then argue for differing approaches toward optimizing choices such as always making choices that minimize regret (ie, integral over time of difference between choice and optimal choice) [20] or enforcing that a choice never be made when a higher-quality one was available [19]. A major limitation of this work is that its baseline standard of fairness is baked into the choice of a quality function, giving a lot of subjective authority to an engineer, and leaving only long-term and short-term optimizations to consider, which can be quite hard to reason about.

3.3 I made X ML system, are the results fair?

This realm of work might be construed as "algorithmic damage control" given that the attempt is not to make a fair algorithm, but rather to develop post-hoc analysis methods that help discover what may be unfair about a black box model. One class of approaches, again premising fairness on prevention of disparate treatment, involves developing mathematical tools to determine whether features related to protected categories (either directly or through co-variables) are influential on the model [2]. A number of these have come in direct relation to the buzz around the problematic recidivism instrument [11, 41]. Given our P1 criterion, these methods, at best, may help an end-user assess whether disparate impact occurred given a set of known protected categories. However, at worst, if we do not know which categories to assess for disparate impact, they may allow a model to pass as fair while unfair biases are still present. In summary, it is a progress that we have ways to verify if an approaches satisfies a disparate impact constraint; however, we are still levied with the challenge of deciding what the constraints should be.

The other major theme in this area of fairness brings back interpretability, but instead of model interpretability, the aim is interpreting why a particular result was obtained [15, 28].

While this has been successful in certain cases where researchers had an intuition about what kind of internal representation the model may have been using, these interpretations rely on naive, simplified approximations of the model. That is, they are unable to interpret the model globally and instead regress on a set of features in a localized subspace. Again, interpretability satisfies certain normative criteria of fairness by giving some power to an end-user to understand a result. However, so long as we cannot interpret the results globally, it's unclear how much power of recourse one may gain from such an interpretation. Further, this approach only allows us to interpret results using a predefined axis, meaning if we must know what we are looking for before this method becomes useful.

4 CONCLUDING REMARKS

In the above, we introduced a robust, normative concept of fairness for ML and evaluated the current literature against it. This evaluation has attempted to shed light on the kinds of authority being wielded when determining "fairness" in an ML system and how that authority operates in different contexts and framings. Before offering summary remarks and recommendations, we would like to make a takeaway point regarding the focus on disparate impact. As has been noted by other researchers [11, 22], disparate impact and making an accurate model are in fundamental tension with one another. Models are discriminators and as such, adding constraints affects our ability to make a good classifier. While, we believe is still useful to characterize fair constraints, it is our belief that this focus on fairness is limited and, regardless of the normative criteria argued for in this paper, will require further approaches to fairness or risk rejection by parts of the community less willing to trade-off accuracy.

With that said, we offer the position that using a normative criterion of fairness (such as P1) should lead the community toward discourse around the legitimacy of authority enacted when deciding what is "fair" in ML. This would support furthering current discussions around the utility of ML interpretation and explanation, enriching our deliberations around how to reason about interpretations and who has the authority to determine when an explanation is fair. Using our criterion, we would suggest that ML practitioners begin incorporating outside experts, review boards, and communities to participate in the discussion of what fair means in different ML problem spaces. Further, we argue there should be elevated attention toward constructing educational content meant for teaching fairness in ML curricula. This content would both catalyze healthy debate within the community and ideally, in the classroom. Finally, we argue that "usable fairness" be taken seriously so that as novel methods are presented we correspondingly value putting those methods to practice and testing them across applications.

REFERENCES

- [1] 2015. Google apologises for Photos app's racist blunder. *BBC News* (July 2015). <http://www.bbc.com/news/technology-33347866>
- [2] Julius Adebayo and Lalana Kagal. 2016. Iterative Orthogonal Feature Projection for Diagnosing Bias in Black-Box Models. *arXiv preprint arXiv:1611.04967* (2016).
- [3] Aws Albarghouthi, Loris D'Antoni, Samuel Drews, and Aditya Nori. 2016. Fairness as a Program Property. *arXiv:1610.06067 [cs]* (Oct. 2016). <http://arxiv.org/abs/1610.06067> arXiv: 1610.06067.
- [4] Mark Andrejevic. 2014. Big Data, Big Questions| The Big Data Divide. *International Journal of Communication* 8, 0 (June 2014), 17. <http://ijoc.org/index.php/ijoc/article/view/2161>
- [5] Igor Bilogrevic and Martin Ortlieb. 2016. "If You Put All The Pieces Together...": Attitudes Towards Data Combination and Sharing Across Services and Companies. ACM Press, 5215–5227. <https://doi.org/10.1145/2858036.2858432>
- [6] Sarah Bird, Solon Barocas, Kate Crawford, Fernando Diaz, and Hanna Wallach. 2016. *Exploring or Exploiting? Social and Ethical Implications of Autonomous Experimentation in AI*. SSRN Scholarly Paper ID 2846909. Social Science Research Network, Rochester, NY. <https://papers.ssrn.com/abstract=2846909>
- [7] Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. *arXiv:1607.06520 [cs, stat]* (July 2016). <http://arxiv.org/abs/1607.06520> arXiv: 1607.06520.
- [8] Simone Browne. 2015. *Dark matters: On the surveillance of blackness*. Duke University Press.
- [9] Joy Buolamwini. 2016. How I'm fighting bias in algorithms. (Nov. 2016). https://www.ted.com/talks/joy_buolamwini_how_i_m_fighting_bias_in_algorithms
- [10] L. Elisa Celis, Amit Deshpande, Tarun Kathuria, and Nisheeth K. Vishnoi. 2016. How to be Fair and Diverse? *arXiv:1610.07183 [cs]* (Oct. 2016). <http://arxiv.org/abs/1610.07183> arXiv: 1610.07183.
- [11] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *arXiv preprint arXiv:1703.00056* (2017).
- [12] Kate Crawford and Jason Schultz. 2013. *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*. SSRN Scholarly Paper ID 2325784. Social Science Research Network, Rochester, NY. <http://papers.ssrn.com/abstract=2325784>
- [13] Amit Datta, Michael Carl Tschantz, and Anupam Datta. 2015. Automated Experiments on Ad Privacy Settings. *Proceedings on Privacy Enhancing Technologies* 2015, 1 (2015), 92–112. <https://doi.org/10.1515/popets-2015-0007>
- [14] Benedict de Spinoza. 1970. *Ethics*. Simon and Schuster.
- [15] Miguel Ferreira, Muhammad Bilal Zafar, and Krishna P Gummadi. 2016. The Case for Temporal Transparency: Detecting Policy Change Events in Black-Box Decision Making Systems. *arXiv preprint arXiv:1610.10064* (2016).
- [16] Alexander R Galloway. 2014. *Laruelle: Against the digital*. University of Minnesota Press Minneapolis.
- [17] Samuel Gibbs. 2015. Musk, Wozniak and Hawking urge ban on warfare AI and autonomous weapons. *The Guardian* (July 2015). <https://www.theguardian.com/technology/2015/jul/27/musk-wozniak-hawking-ban-ai-autonomous-weapons>
- [18] Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. 2016. Fair Learning in Markovian Environments. *arXiv:1611.03071 [cs]* (Nov. 2016). <http://arxiv.org/abs/1611.03071> arXiv: 1611.03071.
- [19] Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. 2016. Fairness in Reinforcement Learning. *arXiv:1611.03071 [cs]* (Nov. 2016). <http://arxiv.org/abs/1611.03071> arXiv: 1611.03071.
- [20] Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. 2016. Rawlsian Fairness for Machine Learning. *arXiv:1610.09559 [cs]* (Oct. 2016). <http://arxiv.org/abs/1610.09559> arXiv: 1610.09559.
- [21] Jeff Larson Lauren Julia Angwin Kirchner, Surya Mattu. 2016. Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And it's Biased Against Blacks. (May 2016). <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [22] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2016. Inherent Trade-Offs in the Fair Determination of Risk Scores. *arXiv:1609.05807 [cs, stat]* (Sept. 2016). <http://arxiv.org/abs/1609.05807> arXiv: 1609.05807.
- [23] Kristian Lum and James Johnrow. 2016. A statistical framework for fair predictive algorithms. *arXiv:1610.08077 [cs, stat]* (Oct. 2016). <http://arxiv.org/abs/1610.08077> arXiv: 1610.08077.
- [24] Lydia Manikonda and Munmun De Choudhury. 2017. Modeling and Understanding Visual Attributes of Mental Health Disclosures in Social Media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 170–181.
- [25] Chantal Mouffe. 1999. Deliberative democracy or agonistic pluralism? *Social Research* (1999), 745–758.
- [26] Frank Pasquale. 2015. *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- [27] Plato and Francis Macdonald Cornford. 1945. *The republic of Plato*. Vol. 30. Oxford University Press London.
- [28] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Why Should I Trust You?: Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1135–1144.
- [29] Richard Rorty. 1989. *Contingency, irony, and solidarity*. Cambridge University Press.
- [30] Richard Rorty. 2009. *Philosophy and the Mirror of Nature*. Princeton University Press.
- [31] Koustuv Saha, Ingmar Weber, Michael L Birnbaum, and Munmun De Choudhury. 2017. Characterizing Awareness of Schizophrenia Among Facebook Users by Leveraging Facebook Advertisement Estimates. *Journal of Medical Internet Research* 19, 5 (2017), e156.
- [32] Olivia Solon. 2017. 'This oversteps a boundary': teenagers perturbed by Facebook surveillance. *The Guardian* (May 2017). <https://www.theguardian.com/technology/2017/may/02/facebook-surveillance-tech-ethics>
- [33] Aaron Springer, Victoria Hollis, and Steven Whittaker. 2017. Dice in the Black Box: User Experiences with an Inscrutable Algorithm. <https://aaai.org/ocs/index.php/SSS/SSS17/paper/view/15372>
- [34] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Queue* 11.3 (2013), 10. <http://dataprivacylab.org/projects/onlineads/1071-1.pdf>
- [35] Craig Timberg. 2016. Racial profiling, by a computer? Police facial-ID tech raises civil rights concerns. (Oct. 2016). https://www.washingtonpost.com/business/economy/face-recognition-tech/2016/10/17/986929ea-41f0-44a2-b2b9-90b495230dce_story.html
- [36] Zeynep Tufekci. 2015. Algorithmic Harms beyond Facebook and Google: Emergent Challenges of Computational Agency. *Colorado Technology Law Journal* 13.2 (2015), 203–218. <http://ctlj.colorado.edu/wp-content/uploads/2015/08/Tufekci-final.pdf>
- [37] Berk Ustun and Cynthia Rudin. 2016. Supersparse linear integer models for optimized medical scoring systems. *Machine Learning* 102, 3 (2016), 349–391.
- [38] Fulton Wang and Cynthia Rudin. 2015. Falling Rule Lists. In *AISTATS*.

- [39] Langdon Winner. 1980. Do artifacts have politics? *Daedalus* (1980), 121–136. <http://www.jstor.org/stable/20024652>
- [40] Hongyu Yang, Cynthia Rudin, and Margo Seltzer. 2016. Scalable Bayesian rule lists. *arXiv preprint arXiv:1602.08610* (2016).
- [41] Ke Yang and Julia Stoyanovich. 2016. Measuring Fairness in Ranked Outputs. *arXiv:1610.08559 [cs]* (Oct. 2016). <http://arxiv.org/abs/1610.08559> arXiv: 1610.08559.
- [42] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P. Gummadi. 2016. Fairness Beyond Disparate Treatment & Disparate Impact: Learning Classification without Disparate Mistreatment. *arXiv:1610.08452 [cs, stat]* (Oct. 2016). <http://arxiv.org/abs/1610.08452> arXiv: 1610.08452.
- [43] Indre Zliobaite. 2015. On the relation between accuracy and fairness in binary classification. *arXiv:1505.05723 [cs]* (May 2015). <http://arxiv.org/abs/1505.05723> arXiv: 1505.05723.