

Fairness Sample Complexity and the Case for Human Intervention

Ananth Balashankar
New York University, Google Research
New York, USA
ananth@nyu.edu

Alyssa Lees
Google Research
New York, USA
alyssalees@google.com

ABSTRACT

With the aim of building machine learning systems that incorporate standards of fairness and accountability, we explore explicit subgroup sample complexity bounds. The work is motivated by the observation that classifier predictions for real world datasets often demonstrate drastically different metrics, such as accuracy, when subdivided by specific *sensitive variable* subgroups. The reasons for these discrepancies are varied and not limited to the influence of mitigating variables, institutional bias, underlying population distributions as well as sampling bias. Among the numerous definitions of fairness that exist, we argue that at a minimum, principled ML practices should ensure that classification predictions are able to mirror the underlying sub-population distributions. However, as the number of sensitive variables increase, populations meeting at the intersectionality of these variables may simply not exist or may not be large enough to provide accurate samples for classification. In these increasingly likely scenarios, we make the case for human intervention and applying situational and individual definitions of fairness. In this paper we present lower bounds of subgroup sample complexity for metric-fair learning based on the theory of Probably Approximately Metric Fair Learning. We demonstrate that for a classifier to approach a definition of fairness in terms of specific sensitive variables, adequate subgroup population samples need to exist and the model dimensionality has to be aligned with subgroup population distributions. In cases where this is not feasible, we propose an approach using individual fairness definitions for achieving alignment. We look at two commonly explored UCI datasets under this lens and suggest human interventions for data collection for specific subgroups to achieve approximate individual fairness for linear hypotheses.

ACM Reference Format:

Ananth Balashankar and Alyssa Lees. 2019. Fairness Sample Complexity and the Case for Human Intervention. In *Proceedings of Where is the Human? Bridging the Gap Between AI and HCI, CHI Workshop (SigChi '19)*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SigChi '19, Glasgow, UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

In recent discussions of ethical ML algorithms, evaluating fairness has been frequently predicated on defining constraints based on specific sensitive variables, such as race or gender. These variables should **not** demonstrate conditionally discriminative behavior while learning classification targets. If care is not taken in the construction of a ML model, works such as [1] and [2] have shown that inequalities in underlying data distributions can be amplified in the predicted output, leading to runaway feedback loops.

Recent works [3] have argued that examining the intersectionality of multiple sensitive variables is crucial for establishing coherent standards of fairness. However, real world data sub-populations often display varying underlying sampling distributions, bias and noise. We argue that principles towards *fair* ML should encourage subgroups to perform individually optimally in a classification task and at a minimum be able to reflect their true underlying population distributions.

As the number of sensitive variables increase, the intersectional subgroup populations tend to decrease in size. In fact, it is not unlikely to have a dataset with a subgroup population of one or zero. In these scenarios, it is evident that any classifier which achieves non-trivial accuracy can never be *fair*. In order to overcome this tradeoff, we look to the rich literature of “individual fairness” which defines fairness with respect to a similarity metric between two individuals and enforces that individuals who are similar are treated similarly, within an error bound [3–5]. We find this definition to be useful in measuring the complexity of a class of machine learning models that attempt to achieve approximate *individual fairness* on a dataset.

Using such a model complexity measure, we can compute the minimum number of samples required in each sub-group to learn approximately “individually fair” models with high probability [6]. Sample complexity provides us a lower bound on the number of samples required from sub-group populations before we can make any claims of learning a *fair* model. This methodology provides us a principled way of thinking about the guarantees provided by any algorithm which claims to learn a fair model, similar to how confidence intervals are used for regression coefficients. If the guarantees aren’t feasible, principled human interventions need to be undertaken to better estimate the confidence of the underlying fair machine learning system.

We have been motivated by the insight that many fairness problems in existing classification tasks for specific subpopulations can be remedied by increased data collection, subject to ethical considerations [7, 8]. As such, we suggest that in the spirit of achieving approximately metric-fair learning, sub-population sample size lower bounds need to be aligned with the actual subgroup

sample sizes in the dataset. In this paper, we highlight the discrepancies in model and sample complexities for sub-populations and the disalignment of the actual sizes with these complexities in two popular UCI datasets. Given that sampling bias leads to unintended consequences in ML classifiers, we show the need to gather new data of specific subgroups in prevalent fairness datasets to avoid perpetrating such conditions induced from biased or insufficient sampling.

2 RELATED WORK

Defining required sample complexity for ML algorithms to achieve low generalization error has been studied extensively in ML theory literature [9, 10]. Probably Approximately Correct (PAC) Learning defines upper bounds on generalization error related to a sample's empirical error and complexity of the hypotheses and datasets [11]. More recently, [12, 13] focused on deriving tighter bounds on deep neural networks in the regime of overparameterization based on compression. Related to this paper's focus on overcoming dimensionality constraints for multiple subgroups, [14] explores the problem of learning a single objective using multiple actors simultaneously, which provides an exponential improvement in sample complexity as opposed to no collaboration. Recent works by [15] also highlight that simple hypotheses cause inequity by design. Our work extends on this, providing a methodology for quantifying sampling bounds for smaller informational datasets and highlights the pitfalls of basing a decision making system on a simple linear hypotheses in some circumstances.

3 LIMITS OF FAIR SENSITIVE VARIABLE SUBGROUP SAMPLING

The use of explicit sensitive variables in real-world scenarios is sometimes a hard constraint. One example is legislation enforcing fairness around disparate impact [16, 17].

In simplified examples, exploring the intersectionality of sensitive variables may be appropriate. For example, in this paper we explore two gender and two race subgroups in the evaluation of the UCI Adult dataset. In our simplified example of subgroups this translates into four separate populations. It is conceivable that in a real-world application, the intersection of gender and race subgroups could extend into **many** different populations. As the intersectionality of subgroups grow, it is likely that a subgroup's sample population will be insufficient. In the case of the UCI German dataset used in this paper, marriage status, with five possible values, is treated as a sensitive variable along with gender. However, in the dataset there were **no** samples containing both the attributes of Female and Single.

It is obvious that there is no model of any complexity that can achieve full fairness with respect to the sensitive variables in this case. Increased data collection is obviously a first step in these situations. However, what if the true subgroup population simply does not exist? In these scenarios, the case for alternate definitions of individual fairness should be explored.

Similarly, in order to satisfy fairness constraints on these increasing intersectionalities of populations, it is necessary that we revisit the framework we use to measure the efficacy of the modeling choices. Despite the impossibility results of achieving fairness

in the extreme case of subgroup sized one, there is still a need to highlight cases where simple (linear) models are inadequately applied in datasets with complex underlying subgroup distributions [18, 19]. The ability to objectively quantify the required model and sample complexity to satisfy these fairness constraints will guide the choices made by practitioners and ML researchers.

4 APPROXIMATE FAIRNESS SAMPLE COMPLEXITY BOUNDS

In order to overcome the impossibility results surrounding individual fairness, the definition of Probably Approximately Metric-Fair learning was first introduced in [6], which formalized the guarantees of fairness expected.

Definition 4.1. Given a data distribution D , hypothesis h trained on a sample S of size m drawn in an i.i.d manner, the hypothesis h is said to be (α, γ) approximately metric-fair with respect to a similarity metric d between two inputs x and x' which discards any variations in sensitive attributes, if

$$Pr_{x, x' \sim D}[|h(x) - h(x')| > d(x, x') + \gamma] \leq \alpha$$

This definition is closely linked with Rademacher model complexity measure [20], which measures the capacity of a class of models (hypotheses) to fit the given dataset with any random error vector for the dataset. We present a relevant result from [6] below which gives us the fairness generalization guarantee for a class of models for a given dataset.

THEOREM 4.1. *If the Rademacher complexity for a hypothesis class H , $R_m(H) = \frac{r}{\sqrt{m}}$ then, for all $\delta, \epsilon_\alpha, \epsilon_\gamma \in (0, 1)$, there exists a sample complexity*

$$m = O\left(\frac{r^2 \cdot \ln(\frac{1}{\delta})}{\epsilon_\alpha^2 \cdot \epsilon_\gamma^2}\right)$$

such that with probability $(1 - \delta)$ over the sample S , $h \in H$, if h is (α, γ) approximately metric-fair on sample S , then h is $(\alpha + \epsilon_\alpha, \gamma + \epsilon_\gamma)$ approximately metric-fair on the underlying distribution D .

This means that if we found a hypothesis h which is probably approximately fair over m samples, then that hypothesis would also be probably approximately fair over the true (unsampled) data distribution, albeit with different parameters.

LEMMA 4.2. *Rademacher complexity bounds for linear hypotheses $H = \{x \rightarrow w \cdot x\}$, is given by $R_m(H) \leq \frac{R \cdot \phi}{\sqrt{m}}$ where $R = \max\|x\|_2$ and $\phi = \max\|w\|_2$ [9]*

Combining Lemma 4.2 with Theorem 4.1, we get that the sample complexity for approximately metric-fairness guarantees for linear hypotheses

$$m = O\left(\frac{R^2 \cdot \phi^2 \cdot \ln(\frac{1}{\delta})}{\epsilon_\alpha^2 \cdot \epsilon_\gamma^2}\right) \quad (1)$$

This complexity applies to all $h \in H$. However, if an exponential-time probably approximate metric-fair learning algorithm, that evaluates empirical error for all hypotheses and selects the h that minimizes it, is applied, it results in a smaller sample complexity

$$m = O\left(\frac{R^2 \cdot \phi^2 \cdot \ln(\frac{1}{\delta})}{\min(\epsilon_\alpha, \epsilon_\gamma)^2}\right) \quad (2)$$

This is equivalent to the information theoretical sample complexity in [6]. Since, the set of linear hypotheses is the one applicable to the widely used linear and logistic regression models applied to the UCI datasets and other crucial interpretable ML problems, this forms the basis of our critical analysis. More complex models like the deep neural networks, are currently out of scope for this analysis, though non-linear models may be more appropriate for the designated task.

5 SUBGROUP SAMPLE COMPLEXITY

In constructing subgroup sample bounds, the above sample complexity measure is applied on various subgroups $\mathcal{P}(G)$ of a dataset. The decision to measure sample complexities per subgroup instead of the overall dataset is motivated by the necessity to highlight the discrepancies in the underlying distributions. Achieving approximate fair learning on the entire dataset is only possible when the sample complexities of the sub-groups are met. If not possible, human interventions in the data collection process by either correcting sampling bias or even collecting new data for subgroups not well represented is advised.

We define membership of an instance to an intersectionality of groups $G = \langle g_1, g_2 \dots g_n \rangle$, for different values of n ($< |X|$) sensitive variables. Calculating the sample complexity for each of the subgroups provides a quantitative measure of the differences between subgroups for a given hypotheses class H . Significant discrepancies indicate differences in the underlying distributions, subgroup sampling methodologies or true labeling functions for these subgroups. In some instances, the sample complexity bound on the overall population may suggest if further sampling should be performed before attempting to learn a single model applied on all subgroups indiscriminately. These complexity metrics provide a clear directive of the order of new samples that needs to be collected before we can make any claims of probably approximately metric-fairness on the hypothesis class. The guarantees associated with it need to match the desired operating surface of $\delta, \epsilon_\alpha, \epsilon_\gamma$ before we can utilize H for safe decision making. Also, since these measures are distribution and hypothesis specific, it provides us a much tighter bound than purely combinatorial measures which work for any distribution and learning algorithms.

5.1 Efficiency bounds

While approximate fair learning ensures that similar data points are classified approximately similarly, there may also be requirements on the minimum accuracy achieved by each of the sub-groups. While this might be at odds with the fairness requirement, it still forms the basis of recent bias mitigating ML literature that rely on constrained optimization [21]. We can obtain a lower bound on the sample complexity to learn a (ϵ, δ) PAC objective function in a collaborative manner [14] among k subgroups. Specifically, there exists a hypothesis h chosen from a class H of VC dimension d , such that it has generalization error less than ϵ on all the k subgroups' distributions $\{D_1, D_2, \dots, D_k\}$ with high probability $1 - \delta$. Current literature also provides algorithms for both a centralized and personalized setting, where either a single function is learned for all k subgroups, or each subgroup can have its own function [14]. In the centralized setting, the sample complexity to learn a PAC algorithm which achieves error of less than ϵ in all k subgroups is

lower bounded by

$$m = O\left(\frac{\ln^2(k)}{\epsilon} \left((d+k) \ln\left(\frac{1}{\epsilon}\right) + k \ln\left(\frac{1}{\delta}\right) \right)\right) \quad (3)$$

For the personalized setting, the sample complexity is $O(\ln(k))$ factor smaller than the centralized sample complexity. For uniform convergence lower bounds where $\epsilon, \delta \in (0, 0.1)$, there exists a PAC learnable hypothesis class of VC dimension d for $m \geq dk(1-\delta)/(4\epsilon)$. This provides a minimum bound for achieving a threshold of error rates across all subgroups, which is compatible with the minimum efficiency fairness intent.

6 EVALUATION

We evaluated the empirical Radmacher and associated sample complexities for two datasets, the UCI Adult Census dataset and the UCI German Credit dataset [22] over a linear hypothesis class.

In the UCI Adult Census dataset, income bracket prediction is based on a set of individual categorical and numerical attributes like age, education, occupation, etc. In this paper, we specify race and gender, as *sensitive variables* that should not inform our model in accurately predicting the income bracket. The computed Radmacher complexities, broken down for the intersection of all subgroup combinations for this dataset shown in Table 2, demonstrate that some subgroups are harder to learn than others for a linear hypothesis.

In order to calculate the discrepancy in sub-group sample complexity for the UCI Adult dataset, approximate metric fairness lower bounds for the underlying distributions were ordered by sub-groups for a given $\delta, \epsilon_\alpha, \epsilon_\gamma$. A condensed table is shown in Table 5. We used one-hot encoding for categorical features and hence the number of dimensions for a linear separating hyperplane d , is 108. With $k = 4$ subgroups, the lower bound for PAC learning is $\frac{432(1-\delta)}{\epsilon}$. To calculate the associated complexity, we trained logistic regression models for each subgroup with 10 fold cross validation to choose regularization weights in [23]. We calculated the maximum norms of the normalized inputs (R) and the coefficients (ϕ). Note that since the lower bound sample complexity for PAC-fair learning is denoted in the big O notation (Equation 1), we can only estimate the ordering or ranking of the sample complexities for the subgroups in Table 5 (higher numbered rank has higher complexity values). The ordering of the actual subgroup sample sizes ($4 > 2 > 3 > 1$) reveals that new samples are needed to match the desired sample complexity ordering ($3 > 2 > 4 > 1$). Specifically, more samples for subgroup 3 need to be gathered than for from subgroups 2 or 4 to ensure the ordering of actual sample sizes align with that of the sample complexities.

In the UCI German Credit dataset, personal financial attributes are present with a target objective of predicting good/bad credit. We treat 'personal status' and 'sex' as sensitive attributes. Since there was no data available for all possible sub-combinations of the 5 personal status options, we treat 'personal status and sex' as a compound feature and examine the 4 subgroups existing with membership in the dataset. The lower bound on the PAC learning sample complexity where $d=61, k=4$ is $\frac{244(1-\delta)}{\epsilon}$. Table 3 and Table 6 show disparity in the order of the actual sample sizes ($3 > 2 > 4 > 1$) as compared to desired sample complexity ($2 > 1 > 3 > 4$). This implies that in the UCI German Credit dataset, more new samples from group 2 than from group 3 should be drawn as prescribed by

Table 1: Discrepancy in Rademacher complexity bounds of linear hypotheses for sensitive subgroups.5
Table 2: UCI Adult dataset

Subgroup (per gender/race)	$R_m(H)$
1 (Female/Black)	0.113
2 (Female/White)	1.002
3 (Male/Black)	2.998
4 (Male/White)	0.649
Overall	0.561

.5
Table 3: UCI German Credit dataset

Subgroup (per status/sex)	$R_m(H)$
1 (Male/Separated)	1.776
2 (Female/Separated-Married)	1.051
3 (Male/Single)	0.349
4 (Male/Married)	0.001
Overall	0.301

Table 4: Comparison of sample complexity ranking for Probably Approximately Metric Fairness with actual subgroup sizes of subgroups1
Table 5: UCI Adult dataset

Subgroup	Sample Complexity Rank	Actual Sample Size (Rank)
1	1	2,129 (1)
2	3	8,642 (3)
3	4	2,616 (2)
4	2	19,174 (4)

1
Table 6: UCI German Credit dataset

Subgroup	Sample Complexity Rank	Actual Sample Size (Rank)
1	3	50 (1)
2	4	310 (3)
3	2	548 (4)
4	1	92 (2)

the sample complexity. Similarly, more samples from subgroup 1 need to be collected than from subgroup 4 in order to remove any inversion in ranking of complexities and actual subgroup sizes to ensure guarantees of probably approximate metric-fair learning using linear hypotheses.

7 CONCLUSION

We have suggested an empirical methodology to quantify the discrepancy between distributions of sensitive subgroups based on the theory of sample complexity for probably approximately metric-fair learning. The evaluation on two extensively used datasets in the fairness ML literature, highlight the shortcomings of claims that a linear hypotheses can be probably approximately fair for these two populations. We demonstrate the need for further sampling from particular subgroups before such a conclusion can be made. Future work to extend this analysis to more complex ML models

may provide a principled standard for ensuring subpopulation fairness. We argue in the development of an ethical AI framework for policy and decision makers, sufficient subgroup sampling and the corresponding model complexity should be of prime focus. When adequate subpopulation samples are not feasible, human intervention is advisable.

REFERENCES

- [1] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *EMNLP*, 2017.
- [2] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Eduardo Scheidegger, and Suresh Venkatasubramanian. Runaway feedback loops in predictive policing. *CoRR*, abs/1706.09847, 2017.
- [3] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *ICML*, 2018.
- [4] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard S. Zemel. Fairness through awareness. In *ITCS*, 2012.
- [5] Cynthia Dwork and Christina Ilvento. Fairness under composition. In *ITCS*, 2019.
- [6] Guy N. Rothblum and Gal Yona. Probably approximately metric-fair learning. *CoRR*, abs/1803.03242, 2018.
- [7] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *FAT*, 2018.
- [8] Irene Chen, Fredrik D. Johansson, and David A Sontag. Why is my classifier discriminatory? *CoRR*, abs/1805.12002, 2018.
- [9] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. Adaptive computation and machine learning. MIT Press, 2012.
- [10] John Shawe-Taylor, Martin Anthony, and Norman Biggs. Bounding sample size with the vapnik-chervonenkis dimension. 42:65–73, 02 1993.
- [11] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, November 1984.
- [12] Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. Towards understanding the role of over-parametrization in generalization of neural networks. *CoRR*, abs/1805.12076, 2018.
- [13] Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. *CoRR*, abs/1802.05296, 2018.
- [14] Avrim Blum, Nika Haghtalab, Ariel D Procaccia, and Mingda Qiao. Collaborative pac learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 2392–2401. Curran Associates, Inc., 2017.
- [15] Jon M. Kleinberg and Sendhil Mullainathan. Simplicity creates inequity: Implications for fairness, stereotypes, and interpretability. 2018.
- [16] Supreme Court of the United States. *Griggs v. duke power co.* 1971.
- [17] 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*, KDD '15, pages 259–268, New York, NY, USA, 2015. ACM.
- [18] Surya MaÅŁ Julia Angwin, JeÅŁ Larson and Lauren Kirchner. How we analyzed the compas recidivism algorithm. 2016.
- [19] Alexandra Chouldechova and Max G'Sell. Fairer and more accurate, but for whom? 06 2017.
- [20] Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *J. Mach. Learn. Res.*, 3:463–482, March 2003.
- [21] Aditya Krishna Menon and Robert C Williamson. The cost of fairness in binary classification. In Sorelle A. Friedler and Christo Wilson, editors, *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pages 107–118, New York, NY, USA, 23–24 Feb 2018. PMLR.
- [22] Dua Dheeru and Efi Karra Taniskidou. UCI machine learning repository, 2017.
- [23] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.