

Crowdsourcing with Fairness, Diversity and Budget Constraints

Naman Goel

naman.goel@epfl.ch

Artificial Intelligence Lab, EPFL

Lausanne, Switzerland

Boi Faltings

boi.faltings@epfl.ch

Artificial Intelligence Lab, EPFL

Lausanne, Switzerland

ABSTRACT

Recent studies have shown that the labels collected from crowdworkers can be discriminatory with respect to sensitive attributes such as gender and race. This raises questions about the suitability of using crowdsourced data for further use, such as for training machine learning algorithms. In this work, we address the problem of fair and diverse data collection from a crowd under budget constraints. We propose a novel algorithm which maximizes the expected accuracy of the collected data, while ensuring that the errors satisfy desired notions of fairness. We provide guarantees on the performance of our algorithm and show that the algorithm performs well in practice through experiments on a real dataset.

CCS CONCEPTS

• **Information systems** → **Crowdsourcing**; *Social networks*; Incentive schemes.

KEYWORDS

Crowdsourcing, Data Quality, Bias, Fairness

ACM Reference Format:

Naman Goel and Boi Faltings. 2019. Crowdsourcing with Fairness, Diversity and Budget Constraints. In *AAAI/ACM Conference on AI, Ethics, and Society (AI/ES '19)*, January 27–28, 2019, Honolulu, HI, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3306618.3314282>

1 INTRODUCTION

Algorithmic decision-making is gaining popularity in many diverse application areas of social importance. Examples include criminal recidivism prediction, stop-and-frisk programs, university admissions, bank loan decisions, screening job candidates, fake news control, information filtering (personalization) and search engine rankings etc. Recently, questions were raised about the fairness of these algorithms. An investigation [26] found COMPAS (a popular software used by courts to predict criminal recidivism risk) racially discriminatory. Other software systems have also been found to be biased against people of different races, genders and political views [3, 18, 20, 25]. This has led to a widespread and legitimate concern about the potential negative influence of such systems on the society [2, 29]. One of the main reasons of algorithmic bias is

the bias in the training datasets. In order to achieve algorithmic fairness, the issue of **data fairness** needs to be addressed first. In many interesting cases, data is directly or indirectly influenced by some kind of human feedback. The influence is obvious and direct if human assigned labels are used as a proxy for ground truth labels. However, human feedback can also indirectly influence the so-called “ground-truth” datasets (when the labels are not human assigned but observed in reality). This is because the ground truth labels can only be collected for a finite number of data points and the selection of data points is often influenced by humans. For example, there are no ground truth labels available for recidivism of people who were never released by the judges. In this paper, we focus only on the direct influence of human feedback on data fairness i.e. the case in which humans assign labels for data.

Crowdsourcing is increasingly used to collect training data labels. Inevitably, crowdworkers have different biases, which are then reflected in the labels collected from the workers. A very recent study [8] conducted on Amazon Mechanical Turk showed that the crowdworkers were equally racially biased as COMPAS in predicting recidivism. The difference in false positive rates of crowd predictions for white and black defendants was significant and nearly equal to that of the predictions made by COMPAS. The same was true for false negative rates also. The bias didn’t change much even when the crowdworkers were not explicitly displayed the race of the defendants.

We consider settings similar to [8]. Workers are asked to provide their answers (or labels) about some tasks with unknown ground truth labels. Every task has some non-sensitive details that are shown to the workers and a sensitive attribute (for example, race) that is not explicitly shown. But the sensitive attribute may potentially be correlated with the non-sensitive task details. A worker inspects the tasks assigned to her and submits labels for the tasks. Each task is assumed to have a ground truth label but the workers don’t have any way of accessing the ground truth. They can only use the task details, their prior knowledge and incomplete information from other sources to make an “educated guess” about the ground truth. The examples of such tasks are “Will a defendant with given personal history recidivate within the next two years or not?” or “Will a candidate with given CV be successful in the job applied for?” or “Is given political news item fake?”. The sensitive attributes in these example tasks are race, gender and political group respectively. Every worker charges a fee for answering the assigned tasks. The requester has a budget constraint on the fees that she can pay to the workers. In this paper, we make the following contributions:

- (1) We propose a novel algorithm for assigning tasks to workers, which optimizes the expected accuracy of labels obtained from crowd while ensuring that the collected labels satisfy desired notions of *error fairness*. The algorithm also ensures diversity of responses by limiting the probability of assigning

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 ACM 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.

AI/ES '19, January 27–28, 2019, Honolulu, HI, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6324-2/19/01...\$15.00

<https://doi.org/10.1145/3306618.3314282>

many tasks to a single worker. Our algorithm works even when the values of the sensitive attribute of the tasks are unavailable or can't be used because of ethical/legal reasons.

- (2) With a novel formulation of the task assignment strategy as a probability distribution over the workers, we can cast the optimization problem as a linear program and avoid the use of integer programming or other graph matching algorithms which are popular in the task assignment literature but are harder to solve exactly and analyze. This also makes our algorithm suitable for online settings in which the requester is not aware of the tasks in advance.
- (3) We use a limited number of gold tasks (tasks with known ground-truth answers) for estimating workers' parameters and then optimally assign non-gold tasks to the workers. We provide performance bounds for our algorithm and show empirical performance on a real dataset.

2 RELATED WORK

Empirical Studies : [8] finds racial discrimination in recidivism prediction tasks on Amazon Mechanical Turk (AMT). [23] analyzes linguistic bias in labels collected through GWAP (Games with a Purpose) on AMT. [24] analyzes the linguistic bias in collaboratively produced biographies. [13] finds discrimination in reputation crowdsourcing systems in online marketplaces.

Proposed Solutions : In an independent and pioneering work, [28] considers the problem of fairness in human decision-making tasks like recidivism prediction, without budget and diversity constraints. Criminal cases with *known* race information arrive in batches of *known* sizes and an MDP based maximum weighted matching algorithm assigns each case to *exactly one* human judge such that the overall utility from decisions of releasing or keeping any defendant is maximized, while ensuring *demographic parity* of release decisions across two races. To the best of our knowledge this is the first and the most recent work to consider settings somewhat similar to ours but our work differs from theirs in several ways. We consider general crowdsourcing settings, in which several assumptions from their model don't hold. In particular, they assume that "true" risk scores of individual defendants are known to the human judges and the case assignment algorithm. In general crowdsourcing settings, one can only hope to have an overall label distribution for the population. In fact, finding the label probability for individual tasks is the very objective of crowdsourcing. Further, it is not immediately clear how their work can be extended for other important fairness definitions. In their model, given true risk scores of the defendants, judges only apply different thresholds for black and white defendants to predict recidivism. The threshold parameters alone can't capture unfairness measures such as unequal error rates. Even if one does improvise the model with more parameters, it remains an open question whether the theoretical conjectures made in the paper are still likely. This is because the conjectures assume that every time a judge gives a decision, the model parameters of the judge are updated. This becomes an issue with error rate parameters since the ground truth labels are not revealed for all tasks in crowdsourcing. [21] considers a different but related problem of bias resulting from adaptive data gathering (when the choice of whether to collect more data of a given type depends on the

data already collected) and propose a differentially private data collection process as a solution.

There is also a lot of work on task assignment in crowdsourcing, which doesn't consider fairness. [27] proposes a greedy knapsack approach to satisfy limits on budget and the number of tasks any worker can solve. [15–17] consider task assignment problem when workers arrive online. [4] proposes optimal gold task assignment when workers' diligence change over time.

Beyond data collection, there is also recent work on making algorithms fair and robust to bias in the training data [9, 12, 14, 19, 30] and on correcting bias in training datasets [5, 10]. Correcting bias in a given dataset requires modifying the feature values and/or the labels in the dataset. In this paper, we aim to collect unbiased and high quality dataset to begin with, relaxing the responsibility and the overhead of such post-processing from data users (for example, data scientists and machine learning engineers).

3 MODEL

Let there be a finite set of n workers and a large pool of tasks with unknown ground truth labels. The data requester randomly chooses tasks from the pool one by one and assigns each to one (or more) worker(s). The requester may not have knowledge of all the tasks in the pool (not even the number of tasks in the pool) in advance. A worker i charges a constant amount of fee c_i for every label she provides. The requester has a budget constraint for the maximum *expected* money to be spent on acquiring one label from a worker.

Let Z be a random variable denoting the sensitive attribute and Y denoting the (unknown) ground truth labels of the tasks such that $Z, Y \in \{0, 1\}$. For the tasks attempted by a worker i , let $\hat{Y}_i \in \{0, 1\}$ denote the labels submitted by the worker. We denote the realizations of random variables Z , Y and \hat{Y}_i by lower case letters z , y and \hat{y}_i respectively and will drop the subscripts for brevity when the context is clear. We will use $[n]$ to denote $\{1, 2, \dots, n\}$. The workers are modeled using their accuracy matrices as follows:

Definition 1 (Accuracy Matrices of a Worker). *The accuracy matrices \mathcal{A}_{iz} , $z \in \{0, 1\}$ of a worker i are two 2×2 row stochastic matrices such that, $\forall y, \hat{y}_i \in \{0, 1\}$, the entry $\mathcal{A}_{iz}[y, \hat{y}_i]$ is the probability of the worker's label on a task being \hat{y}_i given that the sensitive attribute of the task is z and the ground truth label is y .*

The two matrices \mathcal{A}_{i0} and \mathcal{A}_{i1} define the accuracy of the worker i for tasks belonging to the two different values of the sensitive attribute. The accuracy matrix model, also known as the Dawid-Skene model [7] in the crowdsourcing literature, is strong enough to capture different errors (for e.g. false positive and false negative rates) that a worker may make for tasks belonging to a given sensitive attribute value. If a worker is unbiased in the sense that her errors don't depend on the value of sensitive attribute of the task, her two accuracy matrices are identical. Note that this model makes an implicit i.i.d. assumption on a worker's answers.

The requester uses a probabilistic policy to assign the tasks to workers and collects the labels from the workers.

Definition 2 (Crowdsourcing Policy). *A crowdsourcing policy is an n -dimensional stochastic vector S , such that an element $S[i]$, $i \in [n]$ is the probability of assigning any task to worker i , regardless of the sensitive attribute value of the task.*

Note that the requester’s policy doesn’t depend on the value of the sensitive attribute of the task. This is an intentional modeling choice to deal with the situations in which the sensitive attribute values of the tasks may not be available. It may be due to missing data, privacy reasons or legal/ethical requirements of not using the sensitive attribute.

For any task, the requester randomly selects one (or more than one) worker(s) with probabilities specified by the crowdsourcing policy vector S and assigns the task to the selected worker(s). The labels collected from the workers are obviously not guaranteed to be error free. We can define the accuracy matrices of the crowdsourcing policy in the same way as we defined the accuracy matrices of workers.

Definition 3 (Accuracy Matrices of a Crowdsourcing Policy). *The accuracy matrices \mathcal{A}_z , $z \in \{0, 1\}$ of a crowdsourcing policy are two 2×2 row stochastic matrices such that, $\forall y, \hat{y} \in \{0, 1\}$, the entry $\mathcal{A}_z[y, \hat{y}]$ is the probability that a crowdsourced label¹ for a task is \hat{y} given that the sensitive attribute of the task is z and the ground truth label is y .*

We use the letter \mathcal{A} to denote accuracy matrices of the crowdsourcing policy as well as that of the workers but readers can differentiate between the two by noting that \mathcal{A} has an additional subscript i when referring to the matrix of a worker i . It is easy to see that we can express the accuracy matrices of a policy in terms of the accuracy matrices of the workers as follows:

$$\mathcal{A}_z = \sum_{i=1}^n S[i] \cdot \mathcal{A}_{iz} \quad , \forall z \in \{0, 1\} \quad (1)$$

The requester is interested in finding a crowdsourcing policy that maximizes the expected accuracy of the collected labels while ensuring that the data is *fair*, *diverse* and is acquired within budget constraints.

Crowd diversity is a subjective property and is generally defined in terms of the demographics of crowdworkers. In this paper, we work with a given set of crowdworkers and can’t control such a measure of diversity. For settings like these, we define diversity as follows:

Definition 4 (β -Diverse Crowdsourcing Policy). *A crowdsourcing policy is called β -diverse if and only if $\forall i \in [n]$, $S[i]$ is upper bounded by β , where β is a diversity parameter such that $0 \leq \beta < 1$.*

This definition limits the influence of individual workers on the overall crowdsourced dataset and aims to distribute the influence across more workers.

Similar to diversity, fairness is also a subjective property. We use some standard definitions of fairness from the machine learning literature [1, 14, 30].

Definition 5 (False Positive Rate Parity). *A crowdsourcing policy, with accuracy matrices \mathcal{A}_0 and \mathcal{A}_1 , is said to satisfy false positive rate parity if and only*

$$\mathcal{A}_0[0, 1] = \mathcal{A}_1[0, 1]$$

¹We note that the accuracy of a crowdsourcing policy can also be defined in terms of aggregated label when multiple labels per task are collected. But such definitions depend on specific label aggregation algorithms used. However, in many cases, it may be sufficient to assume that the accuracy of a policy with aggregated labels is an increasing function of our accuracy, which is a reasonable assumption.

One can similarly define false negative rate parity, which requires $\mathcal{A}_0[1, 0] = \mathcal{A}_1[1, 0]$.

Definition 6 (Error Rate Parity). *A crowdsourcing policy, with accuracy matrices \mathcal{A}_0 and \mathcal{A}_1 , is said to satisfy error rate parity if and only if it satisfies both false positive rate parity and false negative rate parity, i.e.*

$$\mathcal{A}_0 = \mathcal{A}_1$$

It is easy to see that if all workers are unbiased, any crowdsourcing policy satisfies the above fairness definitions and one only needs to select a policy that maximizes accuracy while satisfying budget and diversity constraints. In this paper, we address the general problem scenario (when workers are not necessarily unbiased).

4 FINDING OPTIMAL CROWDSOURCING POLICY

Let’s first assume that the accuracy matrices of all the workers are known and the requester is interested in finding the optimal crowdsourcing policy maximizing the expected accuracy under budget, fairness and diversity constraints. We model this as a constrained optimization problem. The objective function in the minimization problem is the negative of the expected accuracy of the policy variable S :

$$-\mathbb{E}[\mathcal{A}(S)] = - \sum_{z \in \{0,1\}} P(Z = z) \sum_{y \in \{0,1\}} P_z(Y = y) \sum_{i=1}^n S[i] \mathcal{A}_{iz}[y, y] \quad (2)$$

where $P(Z = z)$ is the known prior probability that any random task in the pool will have sensitive attribute value equal to z and $P_z(Y = y)$ is the known prior probability that any random task with sensitive attribute value z in the pool will have a ground truth label equal to y .

Together with the fairness and diversity constraints, we get the following optimization problem:

$$\begin{aligned} \arg \min_S \quad & - \sum_{z \in \{0,1\}} P(Z = z) \sum_{y \in \{0,1\}} P_z(Y = y) \sum_{i=1}^n S[i] \mathcal{A}_{iz}[y, y] \\ \text{subject to} \quad & \sum_{i=1}^n S[i] = 1 \\ & S[i] \geq 0 \quad , \forall i \in [n] \\ & S[i] \leq \beta \quad , \forall i \in [n] \\ & \mathcal{A}_0[0, 1] - \mathcal{A}_1[0, 1] \leq \alpha \\ & -(\mathcal{A}_0[0, 1] - \mathcal{A}_1[0, 1]) \leq \alpha \\ & \sum_{i=1}^n S[i] \cdot c_i \leq C \end{aligned} \quad (3)$$

The first two constraints are due to the fact that the crowdsourcing policy vectors are probabilistic and so, all elements must be positive and sum to 1. The third is the diversity constraint as formalized in Definition 4. The fourth and fifth constraints together are equivalent to $|\mathcal{A}_0[0, 1] - \mathcal{A}_1[0, 1]| \leq \alpha$. For $\alpha = 0$, we get the exact fairness constraint (false positive rate parity) as formalized in Definition 5. Other fairness constraints can also be similarly included. The last constraint is due to the maximum expected budget (C) that can be spent on acquiring one answer from a worker.

4.1 Estimates of Worker Accuracy Matrices

Until now, we assumed that the accuracy matrices of the workers are known. However, in practice, we need to estimate them. As is common in the literature [22], we assume that the requester has some limited number of gold standard tasks. Gold tasks are the tasks for which the requester not only knows the sensitive attribute value z but also the ground truth label y . We use gold tasks to estimate unknown worker accuracy matrices. Estimating all the entries of the worker accuracy matrices requires that every worker answers some gold tasks of each “type” (the type of a task is specified by its ground truth answer and its sensitive attribute value). We assign N_g tasks of every type to each worker to estimate their accuracy matrices. This is explained further in appendix A. Let $\hat{\mathcal{A}}_{iz}$ be the estimate of the worker accuracy matrices \mathcal{A}_{iz} , $\forall z \in \{0, 1\}$. The optimization problem 3 can now be written as follows, by replacing the accuracy matrices with their estimates:

$$\begin{aligned} \arg \min_S \quad & - \sum_{z \in \{0,1\}} P(Z=z) \sum_{y \in \{0,1\}} P_z(Y=y) \sum_{i=1}^n S[i] \hat{\mathcal{A}}_{iz}[y, y] \\ \text{subject to} \quad & \sum_{i=1}^n S[i] = 1 \\ & S[i] \geq 0, \forall i \in [n] \\ & S[i] \leq \beta, \forall i \in [n] \\ & \hat{\mathcal{A}}_0[0, 1] - \hat{\mathcal{A}}_1[0, 1] \leq \alpha \\ & -(\hat{\mathcal{A}}_0[0, 1] - \hat{\mathcal{A}}_1[0, 1]) \leq \alpha \\ & \sum_{i=1}^n S[i] \cdot c_i \leq C \end{aligned} \quad (4)$$

where,

$$\hat{\mathcal{A}}_z = \sum_{i=1}^n S[i] \cdot \hat{\mathcal{A}}_{iz}, \quad \forall z \in \{0, 1\} \quad (5)$$

This is a linear program, which can be exactly solved in polynomial time. In practice, the simplex method [6] can be used to find the optimal solution efficiently with common optimization libraries like IBM CPLEX and SciPy. Depending on the constraints, the cost and the accuracy matrices of workers, it is possible that no feasible solution exists for the optimization problem. In this case, the requester will have no choice but to relax the constraints.

We will now analyze our algorithm theoretically and empirically. Readers can find a summary of steps of our complete crowdsourcing algorithm in appendix B.

5 THEORETICAL ANALYSIS

When worker accuracy matrices are known, our method is guaranteed to provide the optimal solution, satisfying constraints. However, when estimates of the accuracy matrices are used, two interesting questions arise:

- (1) Does the solution of problem 4 (which is optimal and satisfies fairness constraints only according to the estimated accuracy parameters) also satisfy fairness in reality?
- (2) How much does the requester lose in terms of actual expected accuracy of the policy because of using the estimated accuracy parameters in optimization?

Theorem 1. *With probability at least γ , the solution \hat{S} to the optimization problem 4 satisfies*

$$|\mathcal{A}_0[0, 1] - \mathcal{A}_1[0, 1]| \leq \alpha + \delta$$

where

$$\delta = 2 \sqrt{\frac{-\ln(1 - 2\sqrt{\gamma}) + \ln 2}{2N_g}}; \mathcal{A}_z = \sum_{i=1}^n \hat{S}[i] \mathcal{A}_{iz}, \forall z \in \{0, 1\}$$

and N_g is number of gold tasks.

The theorem states that when we use estimates of the worker accuracy matrices instead of the real matrices, the obtained solution \hat{S} doesn't violate the fairness constraints in reality by more than δ , with probability at least γ .

Theorem 2. *Assuming that the optimal solution \hat{S} of problem 4 satisfies fairness constraints of problem 3 and the optimal solution S of problem 3 satisfies fairness constraints of problem 4, then with probability at least γ'*

$$\mathbb{E}[\mathcal{A}(S)] - \mathbb{E}[\mathcal{A}(\hat{S})] \leq 2n\beta \sqrt{\frac{-\ln(1 - 2\sqrt{\gamma'}) + \ln 2}{2N_g}}$$

where

$$\mathbb{E}[\mathcal{A}(S)] = \sum_{z \in \{0,1\}} P(Z=z) \sum_{y \in \{0,1\}} P_z(Y=y) \sum_{i=1}^n S[i] \mathcal{A}_{iz}[y, y],$$

$$\mathbb{E}[\mathcal{A}(\hat{S})] = \sum_{z \in \{0,1\}} P(Z=z) \sum_{y \in \{0,1\}} P_z(Y=y) \sum_{i=1}^n \hat{S}[i] \mathcal{A}_{iz}[y, y]$$

The theorem provides an upper bound on the loss in real expected accuracy of the crowdsourcing policy, when we use the estimated worker matrices instead of the real accuracy matrices for optimization. Note that in both the theorems, the bounds get better with increasing number of gold tasks.

The proofs of the above theorem are not difficult and depend on a simple application of the Hoeffding's inequality.

6 EXPERIMENTAL EVALUATION

6.0.1 Datasets. We use the following datasets in our experiments.

- (1) **Broward County Dataset** [26]: This dataset contains information about 7214 defendants arrested in Broward County, Florida between 2013 and 2014. The information includes race of the defendants among other non-sensitive attributes such as age, prior charges etc. The dataset also contains ground-truth whether the defendants recidivated within 2 years or not. There are 3696 black defendants and 2454 white defendants in the dataset and the base rate of recidivism is 51.43% among the black defendants and 39.36% among the white defendants.
- (2) **Crowd Judgment Dataset**: [8] randomly selected a subset of 1000 defendants from the Broward County dataset and asked 20 random workers on Amazon Mechanical Turk to predict recidivism for each individual. In total, 400 workers participated in their study and each worker submitted answers for 50 different defendants. The dataset contains these answers collected from the crowd.

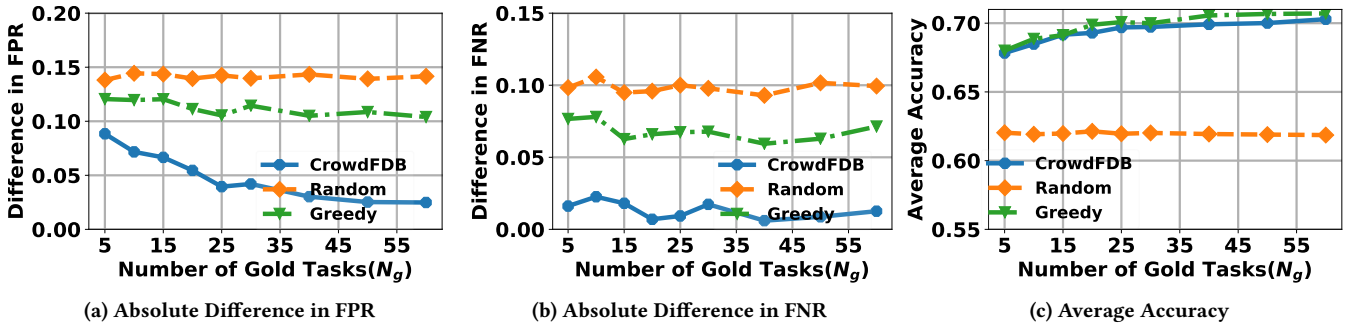


Figure 1: Varying N_g (Number of gold tasks), Settings : Uniform Costs, $\beta = 0.01, \alpha = 0.01$

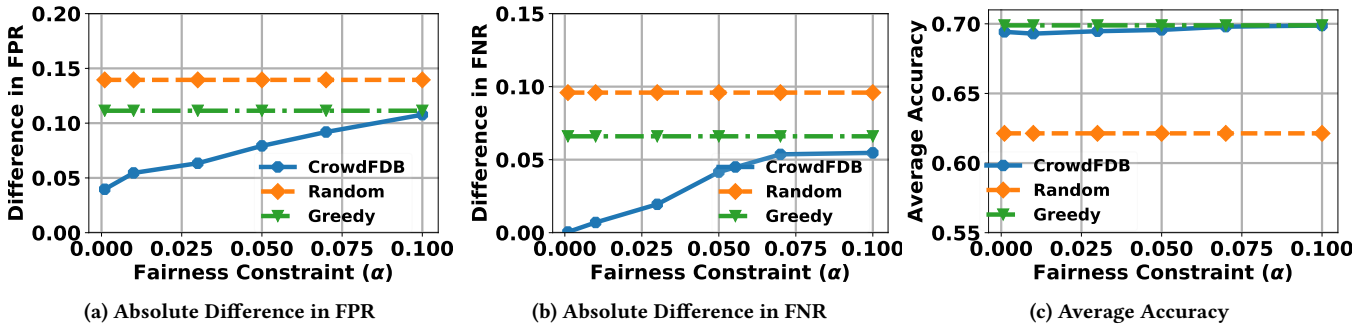


Figure 2: Varying α (Fairness Constraint), Settings : Uniform Costs, $\beta = 0.01, N_g = 20$

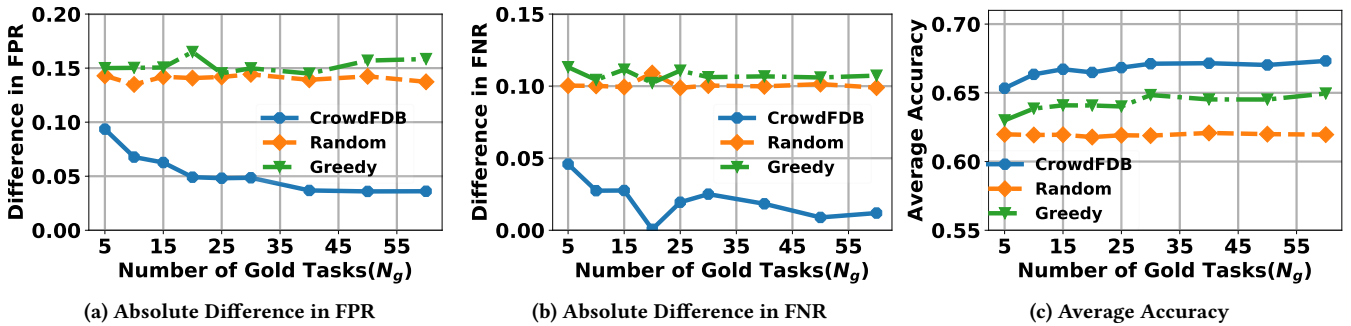


Figure 3: Varying N_g (Number of gold tasks), Settings : Non-Uniform Costs, $\beta = 0.01, \alpha = 0.01, C = 1.5$

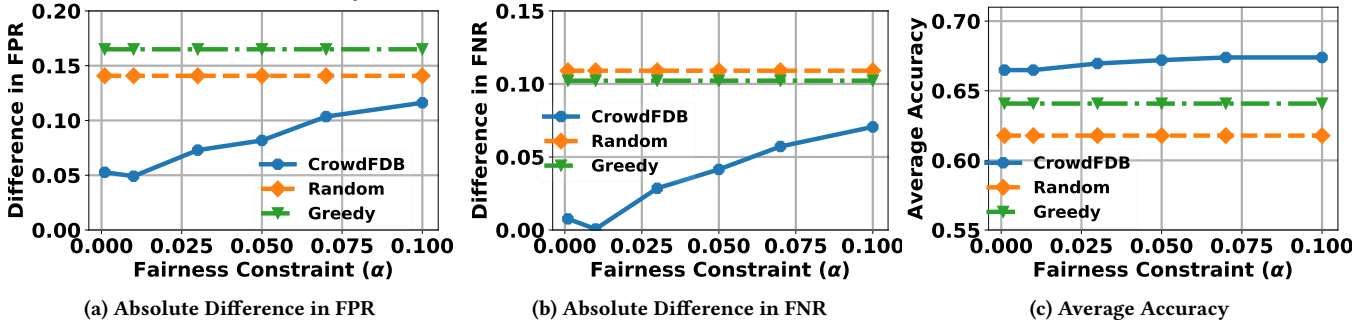


Figure 4: Varying α (Fairness Constraint), Settings : Non-Uniform Costs, $\beta = 0.01, N_g = 20, C = 1.5$

6.0.2 Experiment Outline. The idea is to split the set of defendants into two sets. The first set acts as the gold standard set, which we use to estimate worker accuracy matrices. Once we have the estimates of the worker accuracy matrices, we can solve the optimization problem 4 and learn optimal crowdsourcing policy. We then pick non-gold defendants one by one and assign it to one of the 400 workers, randomly selected according to the policy. The workers' responses are then compared with the ground-truth label to evaluate fairness and accuracy of our crowdsourcing policy.

6.0.3 Handling Limitations of Datasets. Unfortunately, none of the two datasets alone can be used for such experiment. The Broward County dataset contains ground truth labels but doesn't contain workers' answers. On the other hand, the Crowd Judgment dataset does contain worker answers but is very limited for the following reasons. In this dataset, tasks have already been assigned (randomly) to workers and for every defendant, we have responses of only a subset of 20 workers out of all 400 workers. If the crowdsourcing policy learned by our algorithm decides that a worker outside that subset of 20 workers should be assigned a task, then we will need to know the answer of that worker but the answer of this worker is not part of the dataset. The second reason is that every worker has submitted answers for 50 defendants, which is sufficient for getting good estimates of the accuracy parameters of the workers but not big enough to be further split into gold and non-gold sets.

To overcome these limitations, we first create a bigger synthetic dataset using the two real datasets as follows. We generate synthetic answers of all the 400 workers for all the 3696 black and 2454 white defendants in the Broward County dataset. The answers are generated using the worker accuracy parameters estimated from the entire Crowd Judgment dataset. Note that even though this is a synthetic dataset but none of the parameters of the dataset are synthetic. The worker accuracy parameters are derived from the entire real dataset of [8] and the base dataset (Broward County dataset) is used as it is. Indeed, this is not ideal but is perhaps the only option, given the limitations of the available datasets.

Worker Costs : The datasets also don't contain workers' costs. We create this information in two different ways. In the first setting, we associate a uniform cost of \$1 to each worker. In the second and more interesting setting, we probabilistically associate a cost of \$1 or \$3. The probability of a worker's cost being \$3 is equal to her average accuracy and of it being \$1 is equal to 1–her average accuracy. Thus, the higher the average accuracy of a worker, the higher is the probability that she will charge a cost of \$3.

Now this complete dataset is ready to be used in the experiment outlined earlier in this section. We compare our approach (called 'CrowdFDB' in the figures) with two baselines (called 'Random' and 'Greedy' [27]). The baselines are discussed in appendix C.

6.1 Observations

We use equal error rate parity (Definition 6) as the desired fairness. All results reported in the paper are averages over 100 repeated runs. Parameter β was set to 0.01 in the first set of experiments. In the uniform costs settings, C was set to \$1 and in non-uniform settings, $C = \$1.5$.

6.1.1 Uniform Costs. In Figure 1, we keep the fairness constraint α to be fixed (0.01) and observe the effect of increasing number

of gold tasks (N_g). Figures 1a and 1b show that as we increase N_g , the fairness i.e. the absolute difference in FPR (and FNR) for black and white populations, gets closer and closer to α . In other words, the δ of Theorem 1 gets closer to 0 as expected. Moreover, the margin between our algorithm and the baselines also increases. However, meeting the fairness constraints alone is not enough. This could also be done by a bad algorithm that collects equally wrong labels for both white and black populations. Hence, accuracy of the collected labels is also an important measure. Figure 1c shows that our algorithm has an accuracy competitive to the Greedy baseline method, which is a highly efficient baseline in the literature for accuracy optimization. Our algorithm can achieve same level of accuracy while also providing fairness. In Figure 2, we keep N_g fixed (20) and observe the effect of increasing value of α . As value of α increases, the fairness constraints are more relaxed and the algorithm can obtain better accuracy.

6.1.2 Non-Uniform Costs. In the non-uniform costs settings, we observe similar patterns in Figure 3 and 4. There are a few notable differences. The accuracy of our algorithm as well as the Greedy baseline are lower. Our algorithm doesn't select more accurate workers because of budget constraints and the Greedy baseline also finds the density of the more accurate workers comparatively lower due to their high costs and prefers to choose other high density workers. In this case, our algorithm beats Greedy in not just fairness but also in accuracy by better utilizing the available budget.

Figures 5 and 6 of appendix D show the results with $\beta = 0.01, C = 2.5$. Figures 7 and 8 of appendix D show the results with $\beta = 0.005, C = 2.5$. Decreasing the value of parameter β makes the algorithms (ours and the Greedy baseline) more constrained in assigning tasks to the workers that they find to be better. This hits the accuracy of both the algorithms but the general trends discussed above (w.r.t. fairness and accuracy with different N_g and α) remain the same. The effect of increasing budget from 1.5 to 2.5 is that our algorithm can get better accuracy but there is no effect on the performance of other baselines as expected, since there was no explicit budget constraint placed on them.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we addressed the problem of data fairness in crowdsourcing. We proposed a novel crowdsourcing algorithm that learns an optimal selection probability distribution over the available set of workers to maximize the expected accuracy of collected data, while ensuring that the errors in the data are not unfairly discriminatory towards any particular social group. There also remain many challenges to be addressed. These include estimating accuracy without requiring gold standard tasks, relaxing the assumption about knowledge of prior label distribution and the i.i.d. assumption about workers' answers.

Another interesting challenge in slightly different settings is to define data fairness for the case of subjective tasks, which have no ground truth labels and thus, no clear notion of errors. Ensuring fairness in subjective data collection is also likely to create a challenging problem of lying incentives[11] for workers.

REFERENCES

- [1] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2018. Fairness and machine learning : Limitations and Opportunities. <http://fairmlbook.org>. (2018).

- [2] Solon Barocas and Andrew D Selbst. 2016. Big data’s disparate impact. *California Law Review* 671 (2016).
- [3] Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam T Kalai. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in Neural Information Processing Systems*.
- [4] Jonathan Bragg, Daniel S Weld, et al. 2016. Optimal testing for crowd workers. In *International Conference on Autonomous Agents & Multiagent Systems*.
- [5] Flavio Calmon, Dennis Wei, Bhanukiran Vinzamuri, Karthikeyan Natesan Ramamurthy, and Kush R Varshney. 2017. Optimized Pre-Processing for Discrimination Prevention. In *Advances in Neural Information Processing Systems*.
- [6] Vasek Chvatal. 1983. *Linear programming*. Macmillan.
- [7] Alexander Philip Dawid and Allan M Skene. 1979. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Applied statistics* (1979).
- [8] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science advances* (2018).
- [9] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of ITCS*.
- [10] Michael Feldman, Sorelle A Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and removing disparate impact. In *21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [11] Naman Goel and Boi Faltings. 2019. Deep Bayesian Trust: A Dominant and Fair Incentive Mechanism for Crowd. In *AAAI Conference on Artificial Intelligence*.
- [12] Naman Goel, Mohammad Yaghini, and Boi Faltings. 2018. Non-discriminatory machine learning through convex fairness criteria. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- [13] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr.. In *CSCW. 1914–1933*.
- [14] Moritz Hardt, Eric Price, Nati Srebro, et al. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*.
- [15] Chien-Ju Ho, Shahin Jabbari, and Jennifer Wortman Vaughan. 2013. Adaptive task assignment for crowdsourced classification. In *Proceedings of ICML*.
- [16] Chien-Ju Ho and Jennifer Wortman Vaughan. 2012. Online task assignment in crowdsourcing markets. In *Proceedings of AAAI Conference*.
- [17] David R. Karger, Sewoong Oh, and Devavrat Shah. 2014. Budget-Optimal Task Allocation for Reliable Crowdsourcing Systems. *Operations Research* (2014), 1–24.
- [18] Matthew Kay, Cynthia Matuszek, and Sean A Munson. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *33rd Annual ACM Conference on Human Factors in Computing Systems*.
- [19] Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2017. Inherent Trade-Offs in the Fair Determination of Risk Scores. In *Proceedings of ITCS*.
- [20] Juhi Kulshrestha, Motahare Eslami, Johnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. 2017. Quantifying search bias: Investigating sources of bias for political searches in social media. In *ACM Conference on Computer Supported Cooperative Work and Social Computing*.
- [21] Seth Neel and Aaron Roth. 2018. Mitigating Bias in Adaptive Data Gathering via Differential Privacy. In *35th International Conference on Machine Learning*.
- [22] David Oleson, Alexander Sorokin, Greg Laughlin, Vaughn Hester, John Le, and Lukas Biewald. 2011. Programmatic Gold: Targeted and Scalable Quality Assurance in Crowdsourcing. In *AAAI Workshop on Human Computation*.
- [23] Jahna Otterbacher. 2015. Crowdsourcing stereotypes: Linguistic bias in metadata generated via gwap. In *Proceedings of ACM CHI*.
- [24] Jahna Otterbacher. 2015. Linguistic Bias in Collaboratively Produced Biographies: Crowdsourcing Social Stereotypes?. In *Proceedings of AAAI ICWSM*.
- [25] Jahna Otterbacher, Jo Bates, and Paul Clough. 2017. Competent men and warm women: Gender stereotypes and backlash in image search results. In *ACM CHI. ProPublica*. 2017. <https://www.propublica.org/datastore/dataset/compass-recidivism-risk-score-data-and-analysis>. (2017).
- [27] Long Tran-Thanh, Sebastian Stein, Alex Rogers, and Nicholas R Jennings. 2014. Efficient crowdsourcing of unknown experts using bounded multi-armed bandits. *Artificial Intelligence* (2014).
- [28] Isabel Valera, Adish Singla, and Manuel Gomez Rodriguez. 2018. Enhancing the accuracy and fairness of human decision making. In *Advances in Neural Information Processing Systems*. 1774–1783.
- [29] White House. 2016. Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights. *Executive Office of the President* (2016).
- [30] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez Rodriguez, and Krishna P Gummadi. 2017. Fairness Constraints: Mechanisms for Fair Classification. In *20th International Conference on Artificial Intelligence and Statistics (AISTATS)*.

A ESTIMATED ACCURACY MATRICES

Let $\hat{\mathcal{A}}_{iz}$ be the estimate of the worker accuracy matrices \mathcal{A}_{iz} , $\forall z \in \{0, 1\}$. If a worker i answers k tasks correctly out of N_g gold tasks of type $z = 1, y = 1$, then

$$\hat{\mathcal{A}}_{i1}[1, 1] = \frac{k}{N_g} \text{ and } \hat{\mathcal{A}}_{i1}[1, 0] = 1 - \hat{\mathcal{A}}_{i1}[1, 1]$$

Similarly, if she answers k' tasks correctly out of N_g gold tasks of type $z = 1, y = 0$, then

$$\hat{\mathcal{A}}_{i1}[0, 0] = \frac{k'}{N_g} \text{ and } \hat{\mathcal{A}}_{i1}[0, 1] = 1 - \hat{\mathcal{A}}_{i1}[0, 0]$$

This process is repeated with gold tasks of type $z = 0, y = 1$ and $z = 0, y = 0$ to get estimates of all entries of the worker’s matrices.

B SUMMARY OF STEPS OF THE ALGORITHM

Algorithm 1: Crowdsourcing with Fairness, Diversity and Budget Constraints (CrowdFDB)


- 1 Assign N_g gold tasks of every type $((z = 0, y = 0), (z = 0, y = 1), (z = 1, y = 0), (z = 1, y = 1))$ to every worker i in the provided set of n workers.
- 2 Get estimate of every worker i ’s accuracy matrices $\hat{\mathcal{A}}_{i0}$ & $\hat{\mathcal{A}}_{i1}$.
- 3 Solve the linear program to find the best crowdsourcing policy satisfying desired fairness, diversity and budget constraints.
- 4 Pick a task randomly from the task pool of tasks with unknown ground truth labels. Randomly select one or more workers from the set of n workers, with probabilities specified by the crowdsourcing policy. Assign the task to the selected worker(s).
- 5 Repeat Step 4 for all tasks in the pool.

C BASELINES

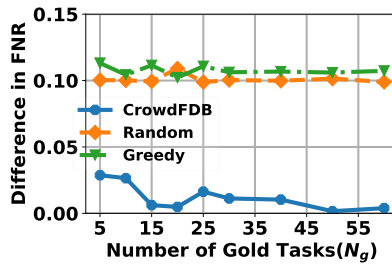
- (1) **Random Policy** : In the random policy, all workers are equally likely to be selected (probability $\frac{1}{400}$) for any task.
- (2) **Greedy Optimization [27]** : In this baseline, we first estimate worker accuracy matrices using gold tasks in exactly same way as we do in our algorithm. However, the optimization is done using a bounded knapsack algorithm instead of linear programming. This algorithm sorts the workers in decreasing order of their “density”, where density is defined as the ratio of the expected accuracy of a worker and her cost. The expected accuracy of a worker can be calculated in the same way as we do for our algorithm i.e. $\sum_{z \in \{0, 1\}} P(Z = z) \sum_{g \in \{0, 1\}} P_z(Y = y) \cdot \hat{\mathcal{A}}_{iz}[y, y]$. Then, the algorithm assigns as many tasks as possible to the highest density worker without violating the diversity constraint. If T is the total number of tasks to be assigned, then a worker can be assigned at most βT tasks. Note that, unlike our algorithm, this baseline has to know the total number of tasks in advance to enforce diversity constraint. Once this worker has reached its capacity, the algorithm starts assigning tasks to the worker with next highest density and continues this for all tasks. However, it may be noted that the Greedy approach was originally proposed for a bit different setting, in which there is an overall crowdsourcing budget and the goal of the requester is to get as many tasks done as possible in that budget, maximizing total utility/accuracy and respecting work limits of the workers.

To give extra advantage to the baselines, we don’t put an explicit budget constraint for them and observe whether they can compete with the fairness and accuracy of our algorithm, which operates under budget constraint.


D ADDITIONAL EXPERIMENTAL RESULTS



(a) Absolute Difference in FPR

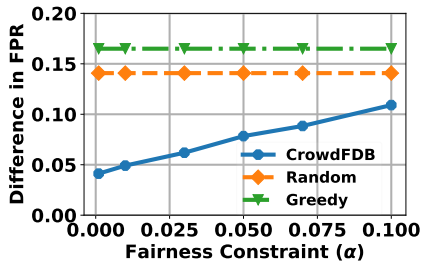


(b) Absolute Difference in FNR

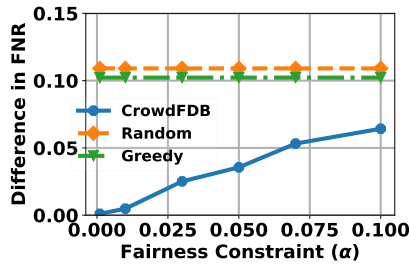


(c) Average Accuracy


Figure 5: Varying N_g (Number of gold tasks), Settings : Non-Uniform Costs, $\beta = 0.01, \alpha = 0.01, C = 2.5$



(a) Absolute Difference in FPR

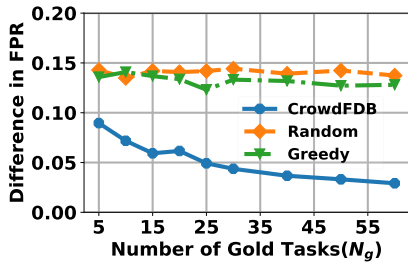


(b) Absolute Difference in FNR

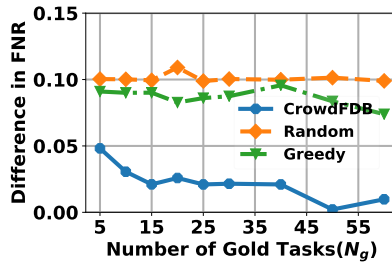


(c) Average Accuracy


Figure 6: Varying α (Fairness Constraint), Settings : Non-Uniform Costs, $\beta = 0.01, N_g = 20, C = 2.5$



(a) Absolute Difference in FPR

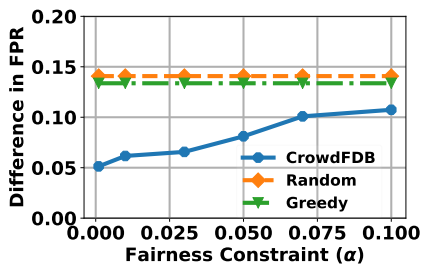


(b) Absolute Difference in FNR

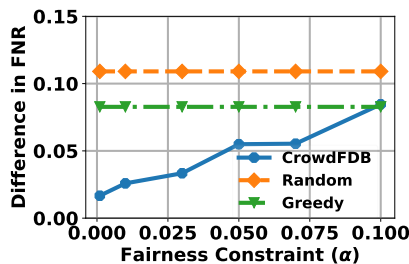


(c) Average Accuracy


Figure 7: Varying N_g (Number of gold tasks), Settings : Non-Uniform Costs, $\beta = 0.005, \alpha = 0.01, C = 2.5$



(a) Absolute Difference in FPR



(b) Absolute Difference in FNR



(c) Average Accuracy

Figure 8: Varying α (Fairness Constraint), Settings : Non-Uniform Costs, $\beta = 0.005, N_g = 20, C = 2.5$