

A Survey on Intersectional Fairness in Machine Learning: Notions, Mitigation, and Challenges

Usman Gohar¹, Lu Cheng²

¹Iowa State University

²University of Illinois Chicago

ugohar@iastate.edu, lucheng@uic.edu

Abstract

The widespread adoption of Machine Learning systems, especially in more decision-critical applications such as criminal sentencing and bank loans, has led to increased concerns about fairness implications. Algorithms and metrics have been developed to mitigate and measure these discriminations. More recently, works have identified a more challenging form of bias called intersectional bias, which encompasses multiple sensitive attributes, such as race and gender, together. In this survey, we review the state-of-the-art in intersectional fairness. We present a taxonomy for intersectional notions of fairness and mitigation. Finally, we identify the key challenges and provide researchers with guidelines for future directions.

1 Introduction

Machine learning (ML) has been increasingly used in high-stake applications such as loans, criminal sentencing, and hiring decisions with reported fairness implications for different demographic groups [Cheng *et al.*, 2021]. Measuring and mitigating discrimination in ML/AI systems has been studied extensively [Mehrabi *et al.*, 2021]. Such works have focused on two specific categories of algorithmic fairness: Group or individual fairness. The majority of early group fairness research was focused on one dimension of group identity, e.g., race or gender. This setting is defined as *independent* groups fairness [Yang *et al.*, 2020a]. However, recent works have identified a more nuanced case of group unfairness that spans multiple subgroups based on Crenshaw’s theory of “intersectionality” [Crenshaw, 1989] called *intersectional group fairness*. At a high level, intersectionality states that interaction along multiple dimensions of identity produces unique and differing levels of discrimination for various possible subgroups, e.g., a Black woman’s experience of discrimination differs from both women and Black people in general. Finally, *gerrymandering* groups are the union of independent and intersectional groups. Figure 1 shows an example of these group fairness definitions using “gender” and “race”.

By categorizing people *only* into distinct overlapping groups, independent group fairness fails to consider the discrimination people face at the intersection of such groups.




Figure 1: Definitions of group fairness [Yang *et al.*, 2020a].

This has been well-studied in philosophy and social psychology (e.g., [Bierly, 1985; Akrami *et al.*, 2011]), but recent works also demand urgency to do so in ML fairness. Specifically, an ML predictor might be fair w.r.t the *independent groups* but not *intersectional groups*. For example, [Buolamwini and Gebru, 2018] identified accuracy disparities that were more significant for Black Women in gender classification algorithms, compared to independent groups. In NLP, works have evaluated popular generative models [Kirk *et al.*, 2021; Tan and Celis, 2019] and also identified such cases of intersectional bias.

Compared to the binary view of fairness in the independent case, the problem of intersectional fairness poses unique challenges. For instance, for what level of granularity of intersectional groups should fairness be guaranteed? On the other hand, smaller subgroups have higher data sparsity, resulting in higher uncertainty [Foulds *et al.*, 2018]. Furthermore, an intersectional identity often amplifies biases that might not exist in its constituent groups (e.g., Black woman vs. Black or Woman), rendering traditional mitigation techniques ineffective. To this end, an emerging body of work, e.g., subgroup fairness [Kearns *et al.*, 2018] and multicalibration [Hebert-Johnson *et al.*, 2018], has proposed various notions of intersectional fairness and mitigation techniques that provide a level of guarantee against intersectional discrimination.

Multiple extensive surveys on fairness in ML have been conducted, such as [Mehrabi *et al.*, 2021] and [Caton and Haas, 2020]. However, they mainly consider the *independent* group fairness and individual fairness while only briefly discussing *intersectional* cases. To bridge this gap, we review the existing fairness literature on intersectional and gerrymandering groups. In particular, we examine existing notions of intersectional fairness in ML and AI and investigate the techniques that enable fair learning for intersectionality.

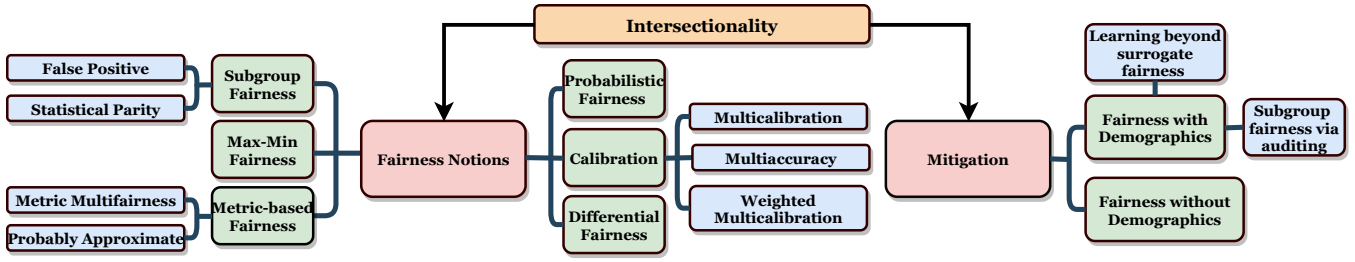


Figure 2: The taxonomy for notions of intersectional fairness and fair learning methods.

Our main contributions are:

1. We propose the first taxonomy (Fig. 2) for the notions of intersectional fairness and fair learning methods for mitigating intersectional discrimination.
2. We thoroughly examine representative intersectional fairness notions and learning methods and discuss their limitations.
3. We conclude with the main challenges faced and point out the open problems in the area.

2 Notions of Intersectional Fairness

Intersectionality, as opposed to group fairness based on independent protected groups (e.g., gender), postulates that the sum of human experiences with discrimination cannot be limited to individual groups alone [Crenshaw, 1989]. Predictors can appear fair when evaluated on independent groups but not at their intersections [Buolamwini and Gebru, 2018]. Satisfying traditional group fairness for intersectionality is infeasible due to potentially infinite overlapping subgroups. This section reviews fairness notions for intersectionality that limit the number of subgroups by balancing the requirements of group fairness and the stronger notion of individual fairness. **Notations.** Each individual is denoted by a tuple (x, y) where $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ denote the instance and ground-truth label, respectively. Let $A = \{s_1, \dots, s_n\}$ be the set of size n protected attributes, f a predictor, and $f(x)$ the predictor output.

2.1 Subgroup Fairness

A pioneering work [Kearns *et al.*, 2018] proposes a stronger notion of group fairness, called *subgroup fairness*, that holds over a large number of *structured* subgroups that can be learned efficiently. In particular, statistical parity (SP) subgroup fairness limits the number of subgroups by disregarding those with limited representations in the data and relaxes the requirement of statistical parity.

Let $\mathcal{C} = \{c : \mathcal{X} \rightarrow \{0, 1\}\}$ be a collection of characteristic functions where $c(s) = 1$ indicates that an individual with protected attribute s is in subgroup c .

Definition 1. $f(x)$ is γ -SP subgroup fair if $\forall c \in \mathcal{C}$:

$$|P(f(x) = 1) - P(f(x) = 1|c(s) = 1)| \times P(c(s) = 1) \leq \gamma, \quad (1)$$

The γ -SP is determined by the worst-case group $c \in \mathcal{C}$. The first term in Eq. 1 is a penalty on the difference in probability between the positive outcome for a specific subgroup c and for the entire population. The smaller the difference, the fairer the outcome. The second term reweighs the difference by the proportion of the size of each subgroup in relation to the population. Consequently, the unfairness of smaller-sized groups is down-weighted in the final γ -SP estimation. Thus, it may not adequately protect small subgroups, even if they have high levels of unfairness. Similarly, subgroup fairness can be applied to the false positive (FP) rate.

2.2 Calibration-based Fairness

Calibration in binary prediction tasks refers to the accuracy of a predictor’s confidence in its predictions [Kearns *et al.*, 2018]. It ensures that the predicted probability distribution for each output class $f(x) = v$ is equal to the actual data probability distribution, i.e., the true expectation is equal to v . For example, if six out of ten samples are positive, the underlying probability and expected predicted probability should also be 0.6. Independently, [Hebert-Johnson *et al.*, 2018] proposes multicalibration that requires all subgroups to be well-calibrated, assuming access to a class of efficiently-learnable characteristic functions. Formally:

Definition 2. Given a parameter $\alpha \in [0, 1]$, $f(x)$ is (\mathcal{C}, α) -multicalibrated if for all predicted values $v \in [0, 1]$, $\forall c \in \mathcal{C}$

$$|\mathbb{E}[c(x) \cdot (y - v) | f(x) = v]| \leq \alpha. \quad (2)$$

The parameter α allows for a less stringent requirement on calibration, i.e., a small miscalibration error α is allowed. Intuitively, a rich class \mathcal{C} will contain groups beyond independent cases, such as intersectional groups, leading to stronger fairness guarantees. *Multiaaccuracy* [Kim *et al.*, 2018] replaces calibration with accuracy constraints to propose a weaker fairness notion, which requires a predictor to be at least α -accurate: $\mathbb{E}[c(x) \cdot f(x) - y(x)] \leq \alpha \forall c \in \mathcal{C}$. Compared to multicalibration, multiaccuracy is less computationally expensive as it is not conditioned on the calibration of each output class across a rich class of intersecting subgroups. These notions define two extremes between efficiency and strong fairness guarantees. To find a balance between the two, [Gopalan *et al.*, 2022] introduces a hierarchy of weighted multicalibration. Formally:

Definition 3. Given \mathcal{C} and a weight class \mathcal{W} , $f(x)$ is $(\mathcal{C}, \mathcal{W}, \alpha)$ -multicalibrated, if $\forall c \in \mathcal{C}$ and $\forall w \in \mathcal{W}$

$$|\mathbb{E}[c(x).w(f(x))(y - f(x))]| \leq \alpha. \quad (3)$$




Figure 3: Hierarchy of multicalibration that interpolates from multi-accuracy (MA) to multicalibration (MC) [Gopalan *et al.*, 2022].

The choice of class \mathcal{W} can lead to multiple variations of the multicalibration notions. *Low-degree multicalibration* is defined as taking the weight function $w(f(x))$ to be a class of $k - 1$ polynomials with k degrees. As shown in Figure 3, when $k = 1$, $w(f(x))$ is constant, and we get the efficient albeit weaker fairness notion of *multiaccuracy*. At higher degrees of polynomial, it converges to *multicalibration*. For a class of 1-Lipstichz functions, we get the (C, α) -*smooth-multicalibration*. Finally, if the predictor is calibrated on prediction intervals instead of calibrating on each predicted value, we arrive at *full-multicalibration*. As such, this hierarchy interpolates the space between multiaccuracy and multicalibration, increasing the strength of fairness guarantees and complexity at higher levels.

2.3 Metric-based Fairness

Another line of work [Yona and Rothblum, 2018] addressed the computational concerns of satisfying fairness for possibly large number of intersectional groups by relaxing the notion of the seminal work of [Dwork *et al.*, 2012] on individual fairness. Individual fairness requires that given a similarity metric if the distance between a pair of individuals is small, a predictor should output similar classification distributions. Inspired by this, [Yona and Rothblum, 2018] proposes a relaxed generalization of individual fairness which allows a small fairness error, called *approximate-metric* fairness. Similar to subgroup fairness and multicalibration works, the relaxation allows the use of efficient learning algorithms that protects every sufficiently-large subgroup. However, unlike those works, the subgroups are not defined a priori. Formally:

Definition 4. For a small $\alpha \in [0, 1]$ and $\gamma \in [0, 1]$, $f(x)$ is (α, γ) -*approximately metric fair* w.r.t similarity metric d and data distribution \mathcal{D} if

$$\mathbb{P}_{(x, x') \sim \mathcal{D}} [|f(x) - f(x')| \geq d(x, x') + \gamma] \leq \alpha, \quad (4)$$

where (x, x') are two individuals sampled from the dataset.

The parameters α and γ allow for small errors in similarity and the metric-fairness measures. Approximate-metric fairness requires that individual fairness holds for all but a small fraction of α pairs of individuals. Consequently, it protects all subgroups of size greater than α as members within the subgroups are treated similarly to those outside. Approximate-metric fairness assumes that the similarity metric is already

known for individuals. To relax the assumption, [Kim *et al.*, 2018] introduces *metric-multifairness*, which supports any similarity metric and requires that similar *subgroups* are treated similarly based on the average distance between individuals in those groups. Formally:

Definition 5. For a small constant $\gamma > 0$ and an unknown similarity metric d , $f(x)$ is (C, d, τ) -*metric multifair* if

$$\mathbb{E}_{(x, x') \sim A} [|f(x) - f(x')|] \leq \mathbb{E}_{(x, x') \sim A} [d(x, x')] + \gamma. \quad (5)$$

More specifically, it requires that individuals in subgroups are treated differently only if they differ substantially from the average difference between individuals within the subgroup. There exist other works (e.g., [Gillen *et al.*, 2018]) that define metric-based fairness in online learning; however, we limit ourselves to the offline setting due to space constraints.

2.4 Differential Fairness

Anti-discrimination laws [Commission, 1978] in the United States declare an outcome as biased if the ratio of probabilities of a favorable outcome between an advantaged and disadvantaged group is less than 0.8. *Differential Fairness* (DF) [Foulds *et al.*, 2020] extends this rule to protect multidimensional intersectional categories. But instead of using a fixed threshold at 80%, DF used a sliding scale, similar to the concept of “differential privacy” [Dwork, 2006], to measure the unfairness of a predictor w.r.t intersectional groups.

Definition 6. $f(x)$ is ϵ -*differentially fair* if

$$e^{-\epsilon} \leq \frac{P(f(x) = y|s_i)}{P(f(x) = y|s_j)} \leq e^{\epsilon}, \quad (6)$$

holds for all tuples $(s_i, s_j) \in A \times A$ where $0 \leq P(s_j) \leq 1$.

For small values of ϵ , the DF criterion states that probabilities of favored outcomes will be similar for any combination of intersectional groups. Unlike other notions, DF ensures fairness for all possible groups, regardless of their size. To estimate the probabilities in Eq. 6, empirical counts for each subgroup can be used. However, it suffers from data sparsity at higher intersections of groups. This can be addressed by using a Dirichlet prior. Finally, [Foulds *et al.*, 2020] also proposes *DF-bias amplification*, which measures the discrimination of a predictor by taking the difference of the DF of the dataset (ϵ_1) and the predictor (ϵ_2).

Furthermore, [Morina *et al.*, 2019] extended the DF notion to other standard group fairness notions such as Statistical Parity, Equality of Opportunity, False Positive Rate Parity, and equalized odds. The ratio in Eq. 6 is replaced with the specific group fairness definition for which the ϵ is measured.

2.5 Max-Min Fairness

The *Max-min* (or *min-max*) notion of fairness is based on the Rawlsian principle of distributive justice [Rawls, 2001]. This principle allows for inequalities but aims to maximize the minimum utility across different protected groups. Given a predictor and a fairness metric, it aims to maximize the fairness of the worst-off subgroup. The Max-Min Fairness [Ghosh *et al.*, 2021] is extended to intersectional cases by measuring the fairness of any combination of intersectional

subgroups using existing fairness definitions and then taking the ratio of the maximum and minimum values from this list of subgroups. A ratio below 1 indicates a disparity between groups, with greater disparity if the ratio is closer to 0. This ratio can be applied to any existing fairness or performance measures like AUC. However, this definition also suffers from low data sparsity when the number of dimensions of intersectionality increases.

2.6 Probabilistic Fairness

Differential fairness uses a Dirichlet prior of uniform parameter α to resolve the issue of intersectional groups having zero counts in the data. The parameter affects this empirical count approach and may miss high-risk subgroups not represented in the data. To solve this, *Probabilistic Fairness* [Molina and Loiseau, 2022] relaxes the requirement of guaranteeing fairness for all subgroups using a probabilistic approach. Formally,

Definition 7. For $\epsilon \geq 0$ and $\delta \in [0, 1]$, a predictor is (ϵ, δ) -probably intersectionally fair if

$$P(U \geq \epsilon) \leq \delta, \quad (7)$$

where $U = u(f(x), s, s')$ measures unfairness for a randomly chosen prediction and two protected groups ($s \neq s'$) to compare them. Probabilistic fairness captures the expected size δ of the population for which the predictor discriminates more than ϵ .

2.7 Discussions

In contrast to traditional fairness metrics, intersectional fairness notions encapsulate a greater number of more granular subgroups and intersectional identities. Existing notions of intersectional fairness depend on the level of unfairness experienced by the most disadvantaged group across all intersections of individual groups. These notions only differ in their approaches for identifying and limiting the vast number of such subgroups to efficiently measure fairness. One risk with designing intersectional notions is that they arbitrarily limit subgroups based on various methods that can efficiently identify them, hence, participating in the same fairness gerrymandering it attempts to solve [Kong, 2022]. For instance, subgroup fairness uses a weight term to prove generalization guarantees w.r.t the underlying population; however, by doing so, it down-weights minority groups which fails to adequately protect them. This highlights the need for a broader involvement of stakeholders to design notions and not simply rely on computational methods. While Max-Min and Differential Fairness truly encapsulates all such groups without disregarding any subgroup, they suffer from data sparsity. Probabilistic approaches might be favorable in such scenarios. Finally, future work could explore the applicability of these definitions in other domains (e.g., recommender systems, NLP, and so on) and define notions for continuous attributes.

3 Improving Intersectional Fairness

ML systems have been shown to exhibit unfairness due to biases in data [Mehrabi *et al.*, 2021] and algorithms [Buolamwini and Gebru, 2018; Gohar *et al.*, 2022]. In response,

there have been great efforts to mitigate bias and improve fairness in ML systems. However, comparatively fewer fair learning algorithms were proposed to address the unique challenges of intersectional fairness. Here, we review the two lines of approaches for intersectional fairness learning: *Intersectional Fairness with Demographics* and *Intersectional Fairness Without Demographics*.

3.1 Intersectional Fairness with Demographics

A surge of methods have been proposed to mitigate bias by learning fair models [Mehrabi *et al.*, 2021]. These techniques are generally applied to the training data (pre-processing), the learning algorithm (in-processing), or the predictions (post-processing). Intersectional fairness with demographics mostly falls into the last two categories using the specific intersectional fairness notions we discussed in Section 2.

Subgroup fairness via auditing. A number of works [Kearns *et al.*, 2018; Kim *et al.*, 2018; Kim *et al.*, 2019; Hebert-Johnson *et al.*, 2018] use auditing to learn fair predictors w.r.t a large number of subgroups. This approach involves an auditor, with access to i.i.d. samples X from an unknown distribution, that assesses the fairness of a predictor using a fairness metric and identifies subgroups with high unfairness. Then a learning algorithm tries to minimize error subject to that fairness constraint. Separately, [Hebert-Johnson *et al.*, 2018; Kearns *et al.*, 2018] both prove that the task of learning such a fair model is equivalent to auditing an arbitrary predictor w.r.t a class of subgroups \mathcal{C} which is computationally equivalent to weak agnostic learning of \mathcal{C} . Utilizing this approach, the seminal work of [Kearns *et al.*, 2018] proposes a zero-sum game between an *Auditor* and a *Learner* for the subgroup fairness notion. In this setting, the zero-sum game is a Fictitious Play using a cost-sensitive classification oracle [Agarwal *et al.*, 2018]. Instead of auditing during training, [Hebert-Johnson *et al.*, 2018] use a post-processing iterative boosting algorithm by combining all $c \in \mathcal{C}$ until the model is α -calibrated. Multiaccuracy [Kim *et al.*, 2019] extends this approach to learning a multi-accurate predictor that guarantees accurate predictions w.r.t \mathcal{C} . Inspired by these approaches, [Kim *et al.*, 2018] propose a variant of stochastic descent gradient that can be leveraged using auditing to post-process a predictor.

Learning beyond surrogate fairness notions. Techniques discussed so far are strictly based on surrogate fairness notions that are adapted to apply to many subgroups e.g., SP-subgroup fairness is a surrogate of the statistical parity notion. Next, we discuss works that go beyond tailor-made intersectional fairness notions. The approach outlined in [Shui *et al.*, 2022] focuses on addressing group sufficiency for many subgroups (including intersectional groups) in ML predictors. Group sufficiency states that given the prediction $f(x)$, the conditional expectation of the ground-truth label ($\mathbb{E}[Y|(f(x), A)]$) is similar across different subgroups. They first derive an upper bound of the group sufficiency gap and propose a bi-level optimization approach with a randomized algorithm that generates the output distribution of the predictor. In the lower level, subgroup-specific output distribution is learned using a small sample of each subgroup’s labeled data.

Then, the final output distribution is updated at the upper level to ensure it is close to all subgroup-specific output distributions. Another work called *GroupFair* [Yang *et al.*, 2020a] proposes Bayes-optimal predictors that are fair for all subgroups w.r.t loss, using a weighted Empirical Risk Minimization (ERM) oracle [Agarwal *et al.*, 2018]. Recently, [Kang *et al.*, 2022] took another step towards a more generalized mitigation approach that does not depend on self-defined fairness notions by capturing linear and non-linear dependence between predictions and intersectional groups using mutual information [Shannon, 1948]. Finally, [Morina *et al.*, 2019], combines and extends previous works [Hardt *et al.*, 2016; Corbett-Davies *et al.*, 2017] to include intersectional cases, by utilizing differential fairness. This involves randomly flipping predictions and a loss function which allows users to find the optimal fairness-accuracy trade-off.

3.2 Intersectional Fairness without Demographics

This line of research addresses intersectional biases without using protected attribute information due to privacy laws. Additionally, data sparsity in smaller subgroups and normative concerns about using synthetic data generation techniques [Wang *et al.*, 2022] make this a compelling approach to tackle intersectional biases. It exploits the correlations between protected attributes and non-protected attributes to approximate the subgroup information. Most of the existing works are in-processing methods.

One such line of work [Hashimoto *et al.*, 2018; Lahoti *et al.*, 2020; Martinez *et al.*, 2021] aims to maximize the minimum utility for all subgroups by using Rawlsian’s Max-Min theory. Unlike parity-based fairness notions, this principle argues in favor of reducing worst-case risk. The fairness objective to minimize the worst-case loss can be formulated as:

$$L_{\max}(f) = \max_{s \in S} \mathbb{E}[l(f; X)]. \quad (8)$$

where $\mathbb{E}[l(f; X)]$ is the expected loss for a loss function e.g., log loss. One approach [Hashimoto *et al.*, 2018] uses *distributionally robust optimization* (DRO) to minimize the worst case loss for any subgroup. DRO achieves this by minimizing the loss over all distributions that are close to the input distribution. Their approach considers the worst-case loss over all distributions with χ^2 -divergence less than r , where r is the radius of a chi-squared ball ($B(P, r)$) around the input probability distribution P . The DRO function is defined as:

$$L_{DRO}(f, r) = \sup_{Q \in B(P, r)} \mathbb{E}_Q[l(f; X)]. \quad (9)$$

Specifically, DRO attempts to reduce the possibly exponential number of subgroups by only considering worst-case distributions that exceed a given size α . A key distinction here is that the objective of the learning algorithm does not depend on α . Consequently, all subgroups have equal representation in the loss function to be minimized. However, this method can potentially optimize noisy outliers, reducing its effectiveness.

To address the limitations of DRO, [Lahoti *et al.*, 2020] propose an Adversarial Reweighting-based approach that relies on the notion of *computationally-identifiable* groups [Hebert-Johnson *et al.*, 2018]. They design a minimax game

between a *learner* and *adversary*: the *learner* is trained to minimize the expected loss while an adversarial neural network is tasked to learn identifiable regions where the learner has significant errors. Their results show that the regions with high errors correspond to various intersectional groups such as *black-female*. The other recent method [Martinez *et al.*, 2021] based on the Max-Min objective proposes a Pareto-efficient [Mas-Colell *et al.*, 1995] learning algorithm to provide a performance guarantee for unidentified protected class w.r.t. to user-defined minimum group size. Although these works target minimizing the worst-case performance of *any* unknown subgroup of a minimum size, experimental results show that they also improve fairness for some intersectional groups.

3.3 Discussion

One of the main drawbacks of current works is that the majority of them rely on specific surrogate fairness notions that we discussed in Section 2. Furthermore, these works ignore certain subgroups that do not conform to specific statistical requirements, e.g., computationally identifiable, that reinforces fairness gerrymandering. One approach to tackle this can be to learn latent representations of intersectional groups that can be then de-biased using geometric approaches [Cheng *et al.*, 2022]. Intersectional fairness without demographics is a promising direction to mitigate intersectional bias, but the current works are limited to the Max-Min notion. While there are certain applications (e.g., healthcare) where improving the utility of the worst-case groups is an important goal, many other applications can be required by law to ensure parity for all subgroups. Similarly, it is critical to evaluate the effectiveness of these methods on different intersectional groups present in the data. For instance, [Lahoti *et al.*, 2020] relies on computational identifiability, which depends on correlations with unprotected attributes. Such methods might fail for intersectional groups that do not have strong demographic signals present in unprotected attributes, hence, failing to protect such groups. Future research can explore learning predictive patterns for underrepresented intersectional groups by leveraging common patterns shared with related groups. For example, Black Females and Black Males might have common structural patterns [Wang *et al.*, 2022].

4 Applications

Most works discussed above are generally focused on classification tasks with i.i.d data. In this section, we review the application of intersectional fairness in other domains of AI.

4.1 Natural Language Processing

Numerous works (e.g. [Tan and Celis, 2019]) have observed that the societal bias inherent in real-world corpora translates to discrimination in NLP models. More recently, there has been a greater effort to focus on benchmarking and debiasing NLP models along intersectional lines.

Benchmark. Several studies have examined bias in sentiment analysis systems, such as [Kiritchenko and Moham-mad, 2018], and found that such systems discriminate based on intersections of gender and race. A closely related study

by [Câmara *et al.*, 2022] across multiple languages confirms these biases. Contextualized word embedding models, including GPT-2 and BERT, have been analyzed for gender and race intersections at sentence level [May *et al.*, 2019] and at contextualized word level [Tan and Celis, 2019]. These works report higher discrimination at the intersection of race and gender (e.g., Black females) compared to either group alone. Separately, [Hassan *et al.*, 2021] evaluates BERT for discrimination against people with disabilities along similar intersectional groups. To automatically identify intersectional biases in static word embeddings, [Guo and Caliskan, 2021] introduces Contextualized Embedding Association Test (CEAT) to measure intersectional bias in contextualized settings. Finally, [Kirk *et al.*, 2021] expands upon these works to incorporate intersections of religion, sexuality, and political affiliations to investigate representational and allocational harms concerning occupational stereotypes in language models. To quantify the scope of the intersectional bias problem in NLP, [Lalor *et al.*, 2022] performs a comprehensive evaluation of state-of-the-art NLP models and debiasing strategies for intersectional bias, benchmarking ten downstream tasks and five demographic groups. These studies highlight the importance of considering a diverse set of intersecting groups in discussions around bias in language models, especially user-facing large language models.

Mitigation. Relatively few works have focused on debiasing along intersectional dimensions. The earliest work by [Subramanian *et al.*, 2021] evaluates two debiasing techniques and shows that debiasing methods based on independent groups are prone to gerrymandering. To address the issue of limited data for intersectional groups, [Cheng *et al.*, 2022] introduces JoSEC, a debiasing approach that leverages the nonlinear geometry of subspace representations to learn intersectional subspace without using predefined word sets. Unlike the linear correlation assumption, they posit that the individual subspaces intersect over a single dimension where the intersectional group subspace resides.

4.2 Ranking Systems

Another common application domain is ranking systems. Fair ranking refers to the method of ensuring that ranking and recommender systems are equitable for all parties involved, including users, providers, and the items being ranked [Mehrabi *et al.*, 2021]. Here we review such works that explore the problem through an intersectional lens.

Top- K Ranking. In the context of fair top- k selection, [Barnabò *et al.*, 2020] examines discrimination along twelve intersections of socioeconomic status, high-school type, and zip code regions for college admissions and proposes an algorithm to select candidates with high utility whilst giving more representation to disadvantaged intersectional groups. Another promising approach [Yang *et al.*, 2020b] uses a causal framework for fair ranking across intersections of gender and race. They compute model-based counterfactuals and rank the resulting scores accordingly. Counterfactual fairness denotes that a prediction is fair if the outcome of an AI system does not change when a single variable is changed and all else remains the same [Kusner *et al.*, 2017].

Fair Rank Aggregation. A similar problem of fair rank aggregation requires combining various rankings to create a consensus ranking, but this can be biased against individual protected attributes like gender, race, and intersectional groups. To resolve this, [Cachel *et al.*, 2022] proposes a group fairness criterion for consensus ranking that ensures fairness for individual groups and their intersections. The unified fairness notion ensures minimal statistical parity difference between pairs of candidate rankings for individual and intersectional groups: $ARP_{a_k} \leq \Delta (\forall a_k \in A \text{ and } IRP \leq \Delta$, where Δ represents desired closeness to statistical parity (zero ensures parity), ARP and IRP represent rank parity for individual and intersectional groups, respectively. They use empirical counts to measure unfairness for each group using this metric, which is prone to data sparsity.

4.3 Auditing and Visualization

Auditing evaluates the fairness of AI systems after training a predictor. It is useful to detect discrimination against a large number of possibly intersecting subgroups. Auditing can identify such subpopulations and make the model more transparent by highlighting its failures. One such work [Kim *et al.*, 2019] leverages a decision-tree-based auditing model to identify bias against dark-skinned women in an image dataset. Other works such as [Hebert-Johnson *et al.*, 2018] and [Kearns *et al.*, 2018] utilize this approach to train fair predictors w.r.t a large number of subgroups.

Some works have created visualization tools to detect potentially discriminatory data subsets. FairVis, [Cabrera *et al.*, 2019], is a visual analytic tool for experts to utilize their domain knowledge in generating subgroups and augmenting automated detection and mitigation strategies. The tool uses clustering analysis to identify statistically similar subgroups and then computes important features and fairness metrics using entropy. Another such tool [Ghai *et al.*, 2021] identifies intersectional biases encoded in word embeddings. Given a pre-trained word embedding, it computes a bias score (using cosine distance) for each subgroup (e.g., male/female for a binary gender) and predefined word sets. A discriminatory word is considered to be associated with an intersectional group if it strongly associates with each of its individual groups according to the bias score. However, this approach may overlook cases like “Hair Weaves”, which are associated with intersectional groups (Black Female) but not individual subgroups (Black or Female).

A software engineering approach [Jin *et al.*, 2020] seeks to ensure adequate representation for relevant intersectional groups within the dataset using coverage. They define intersectional groups using patterns, e.g., {Gender=Female, Race = Black}. Then it requires that intersectional subgroups have a minimum threshold τ of instances. Finally, [Chung *et al.*, 2020] finds intersectional bias in the dataset by dividing it into more granular groups until a subgroup with significant loss is found.

4.4 Discussion

Current applications of intersectionality have been focused on NLP and Ranking. Some other promising applications include recommender systems [Islam *et al.*, 2019], graph em-

Type	Dataset	Demographics
Tabular	Adult	Gender, Age, Race
	Student	Gender, Age, Alcohol, Relationship
	Law School	Gender, Age, Race, Income
	Compass	Gender, Race
NLP	Psychometrics	Gender, Age, Race, Income, Education
	MTC	Gender, Age, Race
	FIPI	Gender, Age, Race, Income, Education
	MBTI	Gender, Age
Ranking	MEPS	Gender, Race, Age
	MovieLens	Gender, Age, Occupation
Image	CelebA	Gender, Age, Race
	UTKFace	Gender, Age, Ethnicity
	PPB	Gender, Race

Table 1: Summary of popular datasets across different AI domains that contain multiple intersectional groups. Law School [Wightman, 1998] and Compass [Angwin *et al.*, 2016] are also used in Ranking. Adult [Kohavi, 1996], Student [Cortez and Silva, 2008], Psychometrics [Abbasi *et al.*, 2021], Multilingual Twitter Corpus (MTC) [Huang *et al.*, 2020], Five Item Personality Inventory (FIPI) and Myers-Briggs Type Indicator (MBTI) [Gjurković *et al.*, 2021], MovieLens [Harper and Konstan, 2015], MEPS [Cohen *et al.*, 2009], CelebA [Liu *et al.*, 2015], UTKFace [Zhang *et al.*, 2017], PPB [Buolamwini and Gebru, 2018].

beddings [Bose and Hamilton, 2019], computer vision [Steed and Caliskan, 2021], and so on. It is imperative that existing ML systems are holistically evaluated under the intersectional framework to aid in developing inclusive and fair ML systems. In NLP, a limitation of current works is the assumption that demographic information is available. Given increasing regulatory and privacy concerns, more research is needed to understand potential correlations in data that can be leveraged to tackle intersectional biases.

5 Datasets and Evaluation Metrics

Available Datasets. Data scarcity is a big challenge for intersectional fairness as the number of dimensions increase. In Table 1, we summarize some popular datasets with adequate intersectional groups, across different AI domains. We hope our consolidated summary provides researchers with convenient access to datasets with rich subgroup information.

Evaluation Metrics. Most studies use the intersectional notions they define for the learning algorithm as the evaluation metrics. Beyond that, worst-case classification accuracy and AUC are broadly used when demographic information is unavailable. These worst-case metrics have also been adopted in NLP and Image classification tasks.

6 Summary and Open Problems

In this survey, we review recent advances in the fairness of ML systems from an intersectional perspective. Intersectionality poses unique challenges that traditional bias mitigation algorithms and metrics cannot effectively address. We review different definitions of group fairness, present a taxonomy of intersectional fairness notions and mitigation methods, and review the literature on intersectionality in other AI domains.

Next, we briefly discuss open problems and potential future research directions.

Data Sparsity. The lack of representative data for marginalized subgroups is a significant challenge. Alternative approaches that do not rely on demographic information may be employed, but these methods do not guarantee that bias against missing subgroups will be addressed. Therefore, a concerted effort to create more inclusive datasets is needed.

Selecting subgroups. Most works propose fairness notions that guarantee fairness only for a limited number of subgroups that are considered statistically meaningful (computationally feasible). This approach fails to protect minority subgroups that do not conform to these statistical requirements. Relying solely on such computational methods reinforces fairness gerrymandering [Kong, 2022]. Hence, it is crucial to involve diverse stakeholders to ensure that the needs and perspectives of different intersectional groups are met.

Generalized mitigation approaches. Existing works on mitigating intersectional bias propose learning algorithms based on specific surrogate fairness notions. These cannot be generalized to other predictors to be used as plug-in mitigation tools. Learning latent representations for intersectional groups so debiased data can be used with any predictor and for any classification task, is a potential direction.

Intersectional fairness beyond parity. Current research overlooks the under-representation of intersectional groups by solely focusing on achieving parity [Kong, 2022]. While it is useful, unequal distribution may be fairer in certain cases. For instance, equalizing hiring rates cannot fix the under-representation of Black females in tech. Therefore, more research on non-distributive intersectional fairness is needed.

Generating test cases for auditing. Generating test cases to audit predictors for intersectional biases is another important direction. With the added complexity of intersectionality, it would be beneficial to evaluate previous testing tools [Chen *et al.*, 2022] and design new tools to test user-facing models for intersectional bias. This can help identify intersectional subgroups against which the predictor is discriminatory.

Beyond mitigation. To effectively address intersectional bias in ML systems, it’s crucial to understand its propagation throughout the ML development cycle, from data collection to algorithms. Exploring causal approaches to understanding intersectional bias is one such interesting direction.

Evaluating fairness notions. Intersectional notions proposed for handling biases have mostly been explored theoretically, with little evidence of their effectiveness on real-world datasets, especially in evaluating the subgroups they fail to protect. Though there are no simple solutions for dealing with intersectional biases in ML, we must measure and benchmark such biases to tackle this problem effectively.

Acknowledgements

This material is based upon work supported by the Cisco Research Gift Grant.

References

- [Abbasi *et al.*, 2021] Ahmed Abbasi, David Dobolyi, John P. Lalor, Richard G. Netemeyer, Kendall Smith, and Yi Yang. Constructing a psychometric testbed for fair natural language processing. In *EMNLP*, pages 3748–3758. ACL, 2021.
- [Agarwal *et al.*, 2018] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. A reductions approach to fair classification. In *ICML*, pages 60–69. PMLR, 2018.
- [Akrami *et al.*, 2011] Nazar Akrami, Bo Ekehammar, and Robin Bergh. Generalized prejudice: Common and specific components. *Psychological Science*, 22(1):57–59, 2011.
- [Angwin *et al.*, 2016] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. In *Ethics of data and analytics*, pages 254–264. Auerbach Publications, 2016.
- [Barnabò *et al.*, 2020] Giorgio Barnabò, Carlos Castillo, Michael Mathioudakis, and Sergio Celis. Intersectional affirmative action policies for top-k candidates selection. *ArXiv*, abs/2007.14775, 2020.
- [Bierly, 1985] Margaret M. Bierly. Prejudice toward contemporary outgroups as a generalized attitude. *Journal of Applied Social Psychology*, 15:189–199, 1985.
- [Bose and Hamilton, 2019] Avishek Bose and William Hamilton. Compositional fairness constraints for graph embeddings. In *ICML*, pages 715–724. PMLR, 2019.
- [Buolamwini and Gebru, 2018] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *In FAT*, pages 77–91. PMLR, 2018.
- [Cabrera *et al.*, 2019] Ángel Alexander Cabrera, Will Epperson, Fred Hohman, Minsuk Kahng, Jamie H. Morgenstern, and Duen Horng Chau. Fairvis: Visual analytics for discovering intersectional bias in machine learning. *IEEE VAST*, 2019.
- [Cachel *et al.*, 2022] Kathleen Cachel, Elke A. Rundensteiner, and Lane Harrison. Mani-rank: Multiple attribute and intersectional group fairness for consensus ranking. *IEEE ICDE*, 2022.
- [Câmara *et al.*, 2022] António Câmara, Nina Taneja, Tamjeed Azad, Emily Allaway, and Richard Zemel. Mapping the multilingual margins: Intersectional biases of sentiment analysis systems in English, Spanish, and Arabic. In *LT-EDI*. ACL, 2022.
- [Caton and Haas, 2020] Simon Caton and Christian Haas. Fairness in machine learning: A survey. *ArXiv*, abs/2010.04053, 2020.
- [Chen *et al.*, 2022] Zhenpeng Chen, J Zhang, Max Hort, Federica Sarro, and Mark Harman. Fairness testing: A comprehensive survey and analysis of trends. *ArXiv*, abs/2207.10223, 2022.
- [Cheng *et al.*, 2021] Lu Cheng, Kush R Varshney, and Huan Liu. Socially responsible ai algorithms: Issues, purposes, and challenges. *JAIR*, 71:1137–1181, 2021.
- [Cheng *et al.*, 2022] Lu Cheng, Nayoung Kim, and Huan Liu. Debiasing word embeddings with nonlinear geometry. In *COLING*, pages 1286–1298, 2022.
- [Chung *et al.*, 2020] Y. Chung, T. Kraska, N. Polyzotis, K. Tae, and S. Whang. Automated data slicing for model validation: A big data - ai integration approach. *IEEE TKDE*, 2020.
- [Cohen *et al.*, 2009] Joel W Cohen, Steven B Cohen, and Jessica S Banthin. The medical expenditure panel survey: a national information resource to support healthcare cost research and inform policy and practice. *Medical care*, pages S44–S50, 2009.
- [Commission, 1978] Equal Employment Opportunity Commission. Guidelines on employee selection procedures. *C.F.R.*, 29., 1978.
- [Corbett-Davies *et al.*, 2017] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. *KDD*, page 797–806. ACM, 2017.
- [Cortez and Silva, 2008] Paulo Cortez and Alice Maria Gonçalves Silva. Using data mining to predict secondary school student performance. *EUROSIS-ETI*, 2008.
- [Crenshaw, 1989] Kimberle Crenshaw. Demarginalizing the intersection of race and sex: A black feminist critique of antidiscrimination doctrine, feminist theory and antiracist politics. *The University of Chicago Legal Forum*, 140:139–167, 1989.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *ITCS*, page 214–226. ACM, 2012.
- [Dwork, 2006] Cynthia Dwork. Differential privacy. In *ICALP 2006*, volume 4052 of *Lecture Notes in Computer Science*, pages 1–12. Springer Verlag, July 2006.
- [Foulds *et al.*, 2018] James R. Foulds, Rashidul Islam, Kamrun Keya, and Shimei Pan. Bayesian modeling of intersectional fairness: The variance of bias. In *SDM*, 2018.
- [Foulds *et al.*, 2020] James R. Foulds, Rashidul Islam, Kamrun N-her Keya, and Shimei Pan. An intersectional definition of fairness. In *ICDE*, pages 1918–1921, 2020.
- [Ghai *et al.*, 2021] Bhavya Ghai, Md Naimul Hoque, and Klaus Mueller. Wordbias: An interactive visual tool for discovering intersectional biases encoded in word embeddings. In *Extended Abstracts of CHI*, pages 1–7, 2021.
- [Ghosh *et al.*, 2021] A. Ghosh, Lea Genuit, and Mary Reagan. Characterizing intersectional group fairness with worst-case comparisons. In *AIDBEI*, 2021.
- [Gillen *et al.*, 2018] Stephen Gillen, Christopher Jung, Michael Kearns, and Aaron Roth. Online learning with an unknown fairness metric. In *NeurIPS*, page 2605–2614, 2018.
- [Gjurković *et al.*, 2021] Matej Gjurković, Mladen Karan, Iva Vukojević, Mihaela Bošnjak, and Jan Snajder. PANDORA talks: Personality and demographics on Reddit. In *SocialNLP*, pages 138–152. ACL, 2021.
- [Gohar *et al.*, 2022] Usman Gohar, Sumon Biswas, and Hridesh Rajan. Towards understanding fairness and its composition in ensemble machine learning. *arXiv preprint arXiv:2212.04593*, 2022.
- [Gopalan *et al.*, 2022] Parikshit Gopalan, Michael P Kim, Mihir A Singhal, and Shengjia Zhao. Low-degree multicalibration. In *COLT*, volume 178 of *PMLR*, pages 3193–3234, 2022.
- [Guo and Caliskan, 2021] Wei Guo and Aylin Caliskan. Detecting emergent intersectional biases: Contextualized word embeddings contain a distribution of human-like biases. In *AIES*, page 122–133. ACM, 2021.
- [Hardt *et al.*, 2016] Moritz Hardt, Eric Price, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *NeurIPS*, volume 29, 2016.
- [Harper and Konstan, 2015] F. Maxwell Harper and Joseph A. Konstan. The movielens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5(4), 2015.
- [Hashimoto *et al.*, 2018] Tatsunori B. Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. Fairness without demographics in repeated loss minimization. In *ICML*, 2018.

- [Hassan *et al.*, 2021] Saad Hassan, Matt Huenerfauth, and Cecilia Ovesdotter Alm. Unpacking the interdependent systems of discrimination: Ableist bias in NLP systems through an intersectional lens. In *EMNLP*, pages 3116–3123. ACL, 2021.
- [Hebert-Johnson *et al.*, 2018] Ursula Hebert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. Multicalibration: Calibration for the (Computationally-identifiable) masses. In *ICML*, volume 80, pages 1939–1948. PMLR, 2018.
- [Huang *et al.*, 2020] Xiaolei Huang, Linzi Xing, Franck Dernoncourt, and Michael J. Paul. Multilingual Twitter corpus and baselines for evaluating demographic bias in hate speech recognition. In *LREC*, pages 1440–1448. ELRA, 2020.
- [Islam *et al.*, 2019] Rashidul Islam, Kamrun Naher Keya, Shimei Pan, and James Foulds. Mitigating demographic biases in social media-based recommender systems. *KDD (Social Impact Track)*, 2019.
- [Jin *et al.*, 2020] Zhongjun Jin, Mengjing Xu, Chenkai Sun, Abolfazl Asudeh, and H. V. Jagadish. Mithracoverage: A system for investigating population bias for intersectional fairness. In *ICDM*, page 2721–2724. ACM, 2020.
- [Kang *et al.*, 2022] Jian Kang, Tiankai Xie, Xintao Wu, Ross Maciejewski, and Hanghang Tong. Infofair: Information-theoretic intersectional fairness. In *IEEE Big Data*, 2022.
- [Kearns *et al.*, 2018] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *ICML*. PMLR, 2018.
- [Kim *et al.*, 2018] Michael P. Kim, Omer Reingold, and Guy N. Rothblum. Fairness through computationally-bounded awareness. In *NeurIPS*, page 4847–4857, 2018.
- [Kim *et al.*, 2019] Michael P. Kim, Amirata Ghorbani, and James Zou. Multiaccuracy: Black-box post-processing for fairness in classification. In *AIES*, page 247–254. ACM, 2019.
- [Kiritchenko and Mohammad, 2018] Svetlana Kiritchenko and Saif Mohammad. Examining gender and race bias in two hundred sentiment analysis systems. In *SemEval*, pages 43–53, 2018.
- [Kirk *et al.*, 2021] Hannah Rose Kirk, Yennie Jun, Haider Iqbal, Elias Benussi, Filippo Volpin, Frédéric A. Dreyer, Aleksandar Shtedritski, and Yuki M. Asano. Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models. In *NeurIPS*, 2021.
- [Kohavi, 1996] Ron Kohavi. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *SIGKDD*, KDD’96, page 202–207. AAAI Press, 1996.
- [Kong, 2022] Youjin Kong. Are “intersectionally fair” ai algorithms really fair to women of color? a philosophical analysis. *FACCT ’22*, page 485–494. ACM, 2022.
- [Kusner *et al.*, 2017] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. Counterfactual fairness. In *NeurIPS*, 2017.
- [Lahoti *et al.*, 2020] Preethi Lahoti, Alex Beutel, Jilin Chen, Kang Lee, Flavien Prost, Nithum Thain, Xuezhong Wang, and Ed Chi. Fairness without demographics through adversarially reweighted learning. In *NeurIPS*, volume 33, pages 728–740, 2020.
- [Lalor *et al.*, 2022] John Lalor, Yi Yang, Kendall Smith, Nicole Forsgren, and Ahmed Abbasi. Benchmarking intersectional biases in NLP. In *NAACL*, pages 3598–3609. ACL, 2022.
- [Liu *et al.*, 2015] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, December 2015.
- [Martinez *et al.*, 2021] Natalia L Martinez, Martin A Bertran, Afroditi Papadaki, Miguel Rodrigues, and Guillermo Sapiro. Blind pareto fairness and subgroup robustness. In *ICML*, PMLR, pages 7492–7501, 2021.
- [Mas-Colell *et al.*, 1995] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [May *et al.*, 2019] Chandler May, Alex Wang, Shikha Bordia, Samuel Bowman, and Rachel Rudinger. On measuring social biases in sentence encoders. In *2019 NAACL*, 2019.
- [Mehrabi *et al.*, 2021] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Comput. Surv.*, 54(6), 2021.
- [Molina and Loiseau, 2022] Mathieu Molina and Patrick Loiseau. Bounding and approximating intersectional fairness through marginal fairness. *arXiv preprint arXiv:2206.05828*, 2022.
- [Morina *et al.*, 2019] Giulio Morina, Viktoriia Oliynyk, Julian Waton, Ines Marusic, and Konstantinos Georgatzis. Auditing and achieving intersectional fairness in classification problems. *CoRR*, abs/1911.01468, 2019.
- [Rawls, 2001] John Rawls. *Justice as fairness: A restatement*. Harvard University Press, 2001.
- [Shannon, 1948] Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [Shui *et al.*, 2022] Changjian Shui, Gezheng Xu, Qi CHEN, Jiaqi Li, Charles Ling, Tal Arbel, Boyu Wang, and Christian Gagné. On learning fairness and accuracy on multiple subgroups. In *NeurIPS*, 2022.
- [Steed and Caliskan, 2021] Ryan Steed and Aylin Caliskan. Image representations learned with unsupervised pre-training contain human-like biases. *FACCT ’21*, page 701–713. ACM, 2021.
- [Subramanian *et al.*, 2021] Shivashankar Subramanian, Xudong Han, Timothy Baldwin, Trevor Cohn, and Lea Frermann. Evaluating debiasing techniques for intersectional biases. In *EMNLP*, pages 2492–2498. ACL, 2021.
- [Tan and Celis, 2019] Yi Chern Tan and L Elisa Celis. Assessing social and intersectional biases in contextualized word representations. In *NeurIPS*, 32, 2019.
- [Wang *et al.*, 2022] Angelina Wang, Vikram V Ramaswamy, and Olga Russakovsky. Towards intersectionality in machine learning: Including more identities, handling underrepresentation, and performing evaluation. In *FACCT*, page 336–349. ACM, 2022.
- [Wightman, 1998] Linda F. Wightman. *LSAC National Longitudinal Bar Passage Study*. LSAC research report series. Law School Admission Council, 1998.
- [Yang *et al.*, 2020a] Forest Yang, Mouhamadou Cisse, and Sanmi Koyejo. Fairness with overlapping groups; a probabilistic perspective. In *NeurIPS*, volume 33, pages 4067–4078, 2020.
- [Yang *et al.*, 2020b] Ke Yang, Joshua R Loftus, and Julia Stoyanovich. Causal intersectionality for fair ranking. 2020.
- [Yona and Rothblum, 2018] Gal Yona and Guy Rothblum. Probably approximately metric-fair learning. In *ICML*, pages 5680–5688. PMLR, 2018.
- [Zhang *et al.*, 2017] Zhifei Zhang, Yang Song, and Hairong Qi. Age progression/regression by conditional adversarial autoencoder. In *CVPR*, pages 5810–5818, 2017.