

COUPLING FAIRNESS AND PRUNING IN A SINGLE RUN: A BI-LEVEL OPTIMIZATION PERSPECTIVE

Anonymous authors

Paper under double-blind review

ABSTRACT

Deep neural networks have demonstrated remarkable performance in various tasks. With a growing need for sparse deep learning, model compression techniques, especially pruning, have gained significant attention. However, conventional pruning techniques can inadvertently exacerbate algorithmic bias, resulting in unequal predictions. To address this, we define a fair pruning task where a sparse model is derived subject to fairness requirements. In particular, we propose a framework to jointly optimize the pruning mask and weight update processes with fairness constraints. This framework is engineered to compress models that maintain performance while ensuring fairness in a single execution. To this end, we formulate the fair pruning problem as a novel constrained bi-level optimization task and derive efficient and effective solving strategies. We design experiments spanning various datasets and settings to validate our proposed method. Our empirical analysis contrasts our framework with several mainstream pruning strategies, emphasizing our method’s superiority in maintaining model fairness, performance, and efficiency.

1 INTRODUCTION

The great success of Artificial Intelligence applications in diverse domains owes much of its achievement to the remarkable capabilities of Deep Neural Networks (DNNs) (He et al., 2016; Dong et al., 2015). To achieve state-of-the-art performance, DNNs often require an immense number of parameters to capture complex relationships within data, which is both computationally and storage intensive. In the pursuit of high performance and efficiency, various DNN model compression techniques have been developed (Han et al., 2015; He et al., 2017; Guo et al., 2016; Molchanov et al., 2016). Although these methods achieve high performance in efficient ways, researchers find that model compression techniques, like pruning, can introduce or exacerbate societal bias (Hooker et al., 2019; Liebenwein et al., 2021; Iofinova et al., 2023). It is crucial to address these concerns to ensure fairness and trust in deploying compressed models into real worlds.

Addressing demographic disparity challenges in machine learning has received much attention in the recent literature on deep learning (Salvador et al., 2022; Wang et al., 2023; Zhu et al., 2023). The primary strategies include fairness-constrained optimization (Zafar et al., 2017) and adversarial machine learning (Madras et al., 2018). Notably, the majority of the current methodologies predominantly target dense deep neural networks that possess a vast parameter space. With a growing need for sparse deep learning, there is an emerging need to delve into the concept of fair pruning to harmonize the triad of performance, fairness, and efficiency. However, achieving this delicate balance in pruning is far from trivial. Network pruning is designed to identify an optimal mask for weights while ensuring high performance, whereas traditional fairness methods prioritize weight adjustments to reduce bias. This dilemma introduces challenges in simultaneously achieving fair-




Figure 1: The accuracy and bias of different pruning methods and Prune & Fair patterns.

ness and identifying suitable masks, given the intertwined relationship between masks, weights, and fairness constraints.

One intuitive approach to derive a fair and compressed model is by sequentially integrating fair learning and model pruning. Yet, this strategy presents several complications. Pruning a fair model prioritizes weights with the goal of maximizing accuracy, often overlooking fairness requirements. Furthermore, when one seeks to tackle this bias in the pruned model, there’s a noticeable degradation in accuracy. In Fig. 1, we preliminarily evaluate four methods at various sparsity levels: (1) the classic structured pruning method (FPGM) (He et al., 2019), (2) the classic unstructured pruning method (Lottery) (Frankle & Carbin, 2018), (3) building a fair model and subsequently pruning it (F&P), and (4) pruning a model and then fine-tuning it with fair constraints (P&F). The results show that pruning inevitably introduces bias while intuitive methods cannot address this challenge.




Figure 2: Sketch of the proposed framework.

To strike a harmonious balance among performance, efficiency, and fairness, a unified strategy is essential, considering the interactions among masks, weights, and constraints. In this work, We present a novel solution called **Bi-level Fair Pruning (BiFP)** based on bi-level optimization to ensure fairness in both the weights and the mask. Fig. 2 shows the overview of our proposed BiFP framework. We first initialize a mask m_0 and optimize the network weights with fairness constraints. Then, we fix the weights, and the relaxed masks are optimized under the consideration of fairness. By iteratively optimizing weights and masks under fairness constraints, our proposed approach effectively tackles the challenge and endeavors to preserve fairness throughout bi-level optimization processes. By **jointly optimizing the mask and weight** update processes while incorporating fairness constraints, we can strike a balance between preserving model performance and ensuring fairness in the pruning task. The proposed research represents a distinctive contribution towards achieving a harmonious balance between model efficiency and fairness. We summarize our contributions as follows.

- We comprehensively investigate the algorithmic bias issue in neural network pruning. We shed light on this unexplored task by comparing various pruning techniques and their potential to introduce bias.
- We introduce new fairness notions tailored specifically for model pruning tasks. These notions provide a more nuanced understanding of fairness in pruning.
- We pioneer a novel joint optimization framework that simultaneously promotes fairness in the mask and weight learning phases during model pruning. This unified approach ensures fairness and circumvents the need for multiple separate optimization runs.
- We perform extensive experiments that not only underline the effectiveness of our approach but also highlight its superiority over existing techniques in terms of fairness, performance, and efficiency. Through ablation studies, we explicitly demonstrate the indispensable role of both the mask constraint and weight constraint towards achieving fair pruning.

2 RELATED WORKS

2.1 FAIR CLASSIFICATION

Recently, algorithmic bias in machine learning has garnered significant attention from the research community, spawning a proliferation of methods and studies. Existing research tackles this challenge in two aspects: 1) defining and identifying bias in machine learning tasks, 2) developing bias elimination algorithms and solutions for conventional machine learning tasks.

To formally define fairness, various fairness notions have been proposed. The most well-recognized notion is *statistical parity*, which means the proportions of receiving favorable decisions for the protected and non-protected groups should be similar. The quantitative metrics derived from *statistical parity* include *risk difference*, *risk ratio*, *relative change*, and *odds ratio* (Zliobaite, 2017). Regarding bias elimination, existing methods are categorized into pre-processing, in-processing, and post-

processing. **Pre-processing** methods modify the training data to remove the potential prejudice and discrimination before model training. Common pre-processing methods include *Massaging* (Kamiran & Calders, 2009-02), *Reweighting* (Calders et al., 2009), and *Preferential Sampling* (Kamiran & Calders, 2012). The **in-processing** methods (Kamiran et al., 2010; Calders & Verwer, 2010; Zafar et al., 2017; Corbett-Davies et al., 2017) tweak the machine learning algorithms to ensure fair predictions, by adding fairness constraints or regularizers into the objective functions in machine learning tasks. The methods for **post-processing** (Kamiran et al., 2012; Hardt et al., 2016; Awasthi et al., 2020) correct the predictions produced by vanilla machine learning models.

2.2 NETWORK PRUNING

Network pruning is a technique aimed at reducing the size and computational complexity of deep learning networks. It involves selectively removing unnecessary weights or connections from a network while attempting to maintain its performance. Network pruning techniques can be broadly categorized into two main types. **Unstructured pruning** approaches remove certain weights or connections from the network without following any predetermined pattern. The early works (Han et al., 2015; Guo et al., 2016; Molchanov et al., 2016) remove weights from a neural network based on their magnitudes. The following works (Baykal et al., 2018; Lin et al., 2020; Yu et al., 2018; Molchanov et al., 2019) prune weights according to training data to approximate the influence of each parameter. Another group of researchers employs optimization methods to address the pruning problems (Ren et al., 2019; Zhang et al., 2018; 2022; Hoefler et al., 2021; Schwag et al., 2020). **Structured pruning** involves pruning specific structures within the neural network, such as complete neurons, channels, or layers. Structured pruning frequently entails the removal of entire filters or feature maps from convolutional layers, resulting in a model architecture that is both structured and readily deployable. This pruning method has been explored extensively in numerous papers (Liu et al., 2019; Li et al., 2019; He et al., 2019; Dong et al., 2017; Ye et al., 2018; 2020). In filter-wise pruning, filters are pruned by assigning an importance score based on weights (He et al., 2017; 2018) or are informed by data (Yu et al., 2018; Liebenwein et al., 2019).

2.3 FAIRNESS IN PRUNED NEURAL NETWORK

There has been increasing attention given to the possible emergence and exacerbation of biases in compressed sparse models, particularly as a result of model pruning processes. Paganini (2020) investigates the undesirable performance imbalances for a pruning process and provides a Pareto-based framework to insert fairness consideration into pruning processes. In the context of text datasets, Hansen & Sjøgaard (2021) provide an empirical analysis, scrutinizing the fairness of the lottery ticket extraction process. Tran et al. (2022) expands upon these insights by demonstrating how pruning might not only introduce biases but also amplify existing disparities in models. Exploring the domain of vision models, Iofinova et al. (2023) delved into the biases that might emerge in pruned models and proposed specific criteria to ascertain whether bias intensification is a probable outcome post-pruning. Lin et al. (2022) introduced a simple yet effective pruning methodology, termed Fairness-aware GRAdient Pruning mEthod (FairGRAPE), that minimizes the disproportionate impacts of pruning on different sub-groups, by selecting a subset of weights that maintain the fairness among multiple groups. Tang et al. (2023) offer a unique perspective to discern fair and accurate tickets right from randomly initialized weights. Despite the notable progress made in the realm of fair model compression, there is a lack of research that couples fair mask and fair weight learning processes in an efficient and simultaneous way.

3 PRELIMINARIES

In this section, we present foundational concepts to ensure a comprehensive understanding of the methodologies and analyses that follow. We begin by exploring the underlying principles of fair machine learning and then delve into the topic of deep neural network pruning.

3.1 FAIR CLASSIFICATION

Consider a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i, s_i)\}_{i=1}^N$, where $\mathbf{x}_i \in \mathcal{X}$ is an input feature, $\mathbf{y}_i \in \mathcal{Y}$ is a ground truth target and $s_i \in \mathcal{S}$ is a sensitive attribute, one can formulate a classification hypothesis space as

$f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ parameterized by θ . The goal of a classification task is to find an optimal parameter θ^* such that: $\theta^* = \arg \min_\theta \mathbb{L}(f_\theta(\mathbf{x}), \mathbf{y})$ where $\mathbb{L} = \frac{1}{N} \sum_{i=1}^N \ell_c(f_\theta(\mathbf{x}_i), \mathbf{y}_i)$ is the loss function and ℓ_c is a surrogate function, such as a hinge function or a logit function (Bartlett et al., 2006).

To take fairness into consideration, it is conventional to adopt fairness constraints into the classic classification task. Specially, we consider two demographic groups $\mathcal{D}^+ = \{(\mathbf{x}_i, \mathbf{y}_i, s_i) | s_i = s^+\}$ and $\mathcal{D}^- = \{(\mathbf{x}_i, \mathbf{y}_i, s_i) | s_i = s^-\}$ denoting the favorable and unfavorable groups, such as the male and female groups in the income prediction task. Given a performance metric $\mathbb{M}(f, \mathcal{D})$, i.e., accuracy, we say that classifier f_θ is fair if the difference between two groups is minor, i.e.,

$$\mathbb{F}(f_\theta, \mathcal{D}; \mathbb{M}) = |\mathbb{M}(f_\theta, \mathcal{D}^+) - \mathbb{M}(f_\theta, \mathcal{D}^-)| \leq \tau, \quad (1)$$

where τ is the user-defined threshold.

Note that the metric \mathbb{M} is tailored according to the specific requirements, goals, or nuances of the application domain. As Iofinova et al. (2023), it is appropriate to measure the accuracy difference with respect to race and gender in the facial identification task since even a moderate difference in accuracy can lead to discrimination in real-world settings. Thus, this paper adopts the disparity of accuracy as the fairness metric. Thus, the metric \mathbb{M} is defined as:

$$\mathbb{A}(f_\theta, \mathcal{D}) = \frac{\sum_{i=1}^N \mathbb{1}(\mathbf{y}_i = f_\theta(\mathbf{x}_i))}{|\mathcal{D}|}, \quad (2)$$

where $\mathbb{1}$ is an indicator function.

Then, we can formulate the fairness-aware classification problem as follows:

$$\theta^* = \arg \min_{\theta} \mathbb{L}(f_\theta(\mathbf{x}), \mathbf{y}) \quad s.t. \quad \mathbb{F}(f_\theta, \mathcal{D}; \mathbb{A}) \leq \tau \quad (3)$$

3.2 NEURAL NETWORK PRUNING

The primary motivation behind network pruning is that many neural network parameters contribute negligibly to the final prediction accuracy. These redundant or insignificant parameters cause over-parameterization, which raises the amount of computation needed during the inference and training phases. Network pruning seeks to mitigate this issue by identifying and eliminating such parameters, thereby achieving a more compact and optimized model.

Pruning techniques try to find a sparse structure (mask m) of a given model f parameterized by θ and have no harm to its accuracy. Then we can formulate it as follows:

$$m = \arg \min_m \frac{1}{n} \sum_{i=1}^n \ell_c(f_{m \odot \theta}(\mathbf{x}_i), \mathbf{y}_i) \quad s.t. \quad \|m\|_0 \leq k, \quad (4)$$

where k is the desired remaining parameters. m represents a binary mask designed to set specific parameters to a value of 0. $\|\cdot\|_0$ denotes the number of non-zero elements. The operator \odot represents the element-wise product. Following the convention, we use sparsity to describe the ratio of zero parameters in a neural network, which can be computed by $1 - \frac{\|m\|_0}{\|\theta\|_0}$.

4 FAIR NEURAL NETWORK PRUNING

Having established the foundational concepts of both fairness in machine learning and the intricacies of neural network pruning, we now converge on a critical intersection of these domains. The challenge lies in effectively integrating fairness considerations into the pruning process, ensuring that the efficiency gains of model compression do not inadvertently compromise the equitable performance of the model. In this section, we formulate the fair learning framework in the neural network pruning process and introduce a novel method that seamlessly marries these two domains, offering a pathway to achieve both compactness in neural networks and fairness in their predictions.

4.1 FAIRNESS NOTIONS FOR COMPRESSED MODELS

Given any arbitrary compressed model, we aim to figure out whether this model produces fair and equalized results for demographic groups. One can measure the metric difference of the compressed

model on two demographic groups. If the compressed model has no difference regarding the selected metric, this compressed model is considered to achieve fairness in terms of the selected fairness notion. Formally, we use f_c to represent the compressed model and formulate the fairness definitions as follows.

Definition 1 (Performance Fairness) *Given a compressed model f_c , we say the model is fair for performance if and only if its metric difference for two demographic groups is minor, i.e., $\mathbb{F}(f_c, \mathcal{D}; \mathbb{A}) = |\mathbb{A}(f_c, \mathcal{D}_{s^+}) - \mathbb{A}(f_c, \mathcal{D}_{s^-})| \leq \tau$.*

In addition to the absolute measurement for the compressed model, we observe that the compressed model may amplify the bias compared to the original model f . The reason is that the model compression process introduces excess performance decrease for the unfavorable group. To capture the exacerbation bias induced by the model compression, we further introduce a metric for the accuracy shrinkage between the compressed model and the original model.

Definition 2 (Compressed Model Performance Degradation) *Given the original model f and its compressed version f_c , the performance reduction is defined as: $\mathbb{R}(f, f_c, \mathcal{D}) = \mathbb{A}(f, \mathcal{D}) - \mathbb{A}(f_c, \mathcal{D})$.*

Further, we define a new fairness metric regarding the performance reduction of model compression.

Definition 3 (Performance Degradation Fairness) *Given the original model f and its compressed version f_c , we say the models f and f_c are fair for performance degradation if and only if their metric decrease difference for two demographic groups is minor, i.e., $\mathbb{F}(f, f_c, \mathcal{D}; \mathbb{R}) = |\mathbb{R}(f, f_c, \mathcal{D}^+) - \mathbb{R}(f, f_c, \mathcal{D}^-)| \leq \tau$.*

4.2 FAIR AND EFFICIENT LEARNING VIA THE LENS OF BI-LEVEL OPTIMIZATION

Following our exploration of various fairness notions, we pivot toward the learning paradigms employed in the realm of fair and compressed neural networks. The intuitive integration of pruning, which focuses on applying masks to neural network weights, and fairness, which emphasizes adjustments to these weights, might seem trivial. In fact, the simplistic amalgamation might even undermine the objectives of both, as shown in Fig. 1, where both performance and fairness decrease. This drives us to reconsider the conventional methodology by designing a holistic solution. Our proposition is a simultaneous approach that enforces fairness in both masks and weights during the whole pruning phase. To this end, we cast the challenge as a bi-level optimization problem, and in the subsequent sections, we detail an efficient solution to this intricate puzzle.

Formally, neural network pruning aims to find a mask m that determines the sparse pattern of the model, while fairness classification aims to acquire a fair parameter set θ that achieves fairness and maintains accuracy. To ensure the fairness requirement in both pruning masks and weights, we consider the model specified by an arbitrary mask m and weights θ as $f_m : m \odot \theta; \mathcal{X} \rightarrow \mathcal{Y}$. Then, we formulate pruning and fairness classification as the following bi-level optimization problem:

$$\begin{aligned} \min & \frac{1}{n} \sum_{i=1}^n \ell_c(f_{m \odot \theta^*(m)}(\mathbf{x}_i), \mathbf{y}_i), \\ \text{s.t. } & \theta^*(m) = \arg \min_{\theta} \frac{1}{n} \sum_{i=1}^n \ell_c(f_{m \odot \theta}(\mathbf{x}_i), \mathbf{y}_i), \quad \mathbb{F}(f_{m \odot \theta}, \mathcal{D}; \mathbb{M}) \leq \tau \end{aligned} \quad (5)$$

where m and θ are the upper-level and lower-level optimization variables, respectively.

Bi-level optimization problems are sensitive with respect to the addition of constraints. Even adding a constraint that is not active at an optimal solution can drastically change the problem (Macal & Hurter, 1997). The constraint in Eq. 5 involving both the upper and the lower level variables can cause problems, i.e., the upper variable might not be able to accept certain optimal solutions of the lower level variable (Mersha & Dempe, 2006). In order to address this problem, we relax this constraint and discuss an efficient solution in the following subsection.

4.3 CONVEX RELAXATION OF FAIRNESS CONSTRAINT

While we’ve posited a broad fairness constraint tailored for fair pruning, the generality of this constraint allows it to be effectively applied across a multitude of contexts. To shed light on its practical implications and to elucidate its inner workings, we delve into a specific example in this section. We consider the equalized accuracy constraint (Hardt et al., 2016) where the accuracy across different demographic groups is similar. i.e. $\mathbb{A}(f_{m \odot \theta}, \mathcal{D}^+) - \mathbb{A}(f_{m \odot \theta}, \mathcal{D}^-) \leq \tau$. To foster equalized accuracy constraint, we formulate the fairness constraint as follows:

$$\mathbb{F}_{\mathbb{A}} = \mathbb{E}_{X|S=1}[\mathbb{1}_{Y \cdot f_c(X) > 0}] - \mathbb{E}_{X|S=-1}[\mathbb{1}_{Y \cdot f_c(X) > 0}] \quad (6)$$

where $\mathbb{F}_{\mathbb{A}}$ is the abbreviation of $\mathbb{F}(f_{m \odot \theta}, \mathcal{D}; \mathbb{A})$, $\mathbb{1}(\cdot)$ is indicator function, and $f_c := f_{m \odot \theta}$. To ensure smooth and differentiable, we reformulate Eq. 6 as follows:

$$\begin{aligned} \mathbb{F}_{\mathbb{A}} &= \mathbb{E}_{X|S=1}[\mathbb{1}_{Yf(X) > 0}] - \mathbb{E}_{X|S=-1}[\mathbb{1}_{Yf(X) > 0}] \\ &= \mathbb{E}_X \left[\frac{P(S=1|X)}{P(S=1)} \mathbb{1}_{(Yf(X)) > 0} \right] + \mathbb{E}_X \left[\frac{P(S=-1|X)}{P(S=-1)} \mathbb{1}_{(Yf(X)) < 0} \right] - 1 \end{aligned} \quad (7)$$

The indicator function can be further replaced with the differentiable surrogate function $u(\cdot)$:

$$\mathbb{F}_{\mathbb{A}} = \frac{1}{N} \sum_{i=1}^N \frac{u(\mathbf{y}_i \cdot f_{m \odot \theta^*(m)}(\mathbf{x}_i))}{P(S_i)} - 1 \quad (8)$$

Eq. 8 is convex and differentiable if the surrogate function u is convex and differentiable at zero with $u'(0) > 0$ (Wu et al., 2019). In the general cases where u is not convex, the differential nature of Eq. 8 ensures efficient optimization for gradient descent.

4.3.1 UPDATING RULE OF THE INNER FUNCTION.

After obtaining the differentiable variant for the fairness constraint, we refine the updating strategy for gradient descent of the inner function as follows:

$$\theta^t = \theta^{t-1} - \alpha \cdot \nabla_{\theta} \left(\frac{1}{n} \sum_{i=1}^n \ell_c + \mathbb{F}_{\mathbb{A}} \right)$$

Thus, the differentiable constraint can be seamlessly integrated with the inner optimization.

4.3.2 UPDATING RULE OF THE UPPER FUNCTION.

For the upper optimization part, the masking variable is binary and discrete. We follow the conventional practice and relax the binary masking variables to continuous masking scores $m \in [0, 1]$ inspired by (Ramanujan et al., 2020). Then, we derive the gradient updating rule for the mask m in the upper objective function:

$$\text{grad}(m) = \nabla_m \ell[(m \odot \theta^*(m)) + \mathbb{F}_{\mathbb{A}}] + \frac{d(\theta^*(m))}{dm} \nabla_{\theta} \ell[(m \odot \theta^*(m)) + \mathbb{F}_{\mathbb{A}}],$$

where ∇_m and ∇_{θ} denote the partial derivatives of the bi-variate function $\ell(m \odot \theta)$. Then, we can update m^t by following step:

$$m^t = m^{t-1} - \beta \cdot \text{grad}(m) \quad (9)$$

In this section, we define the fair learning task in neural network pruning and refine the bi-level optimization task to achieve fair and efficient pruning. Recognizing the challenges posed by the original formulation, we introduce a relaxation to accommodate convex and differentiable functions, greatly facilitating the optimization process. With these differentiable properties, we update the solving strategies intrinsic to bi-level optimization, ensuring their alignment with our tailored objectives.

5 EXPERIMENT

We evaluate our proposed method with several baselines in this section. The models are all implemented using PyTorch and evaluated with Intel(R) Core(TM) i9-10900X CPU and NVIDIA GeForce RTX 3070 GPU.

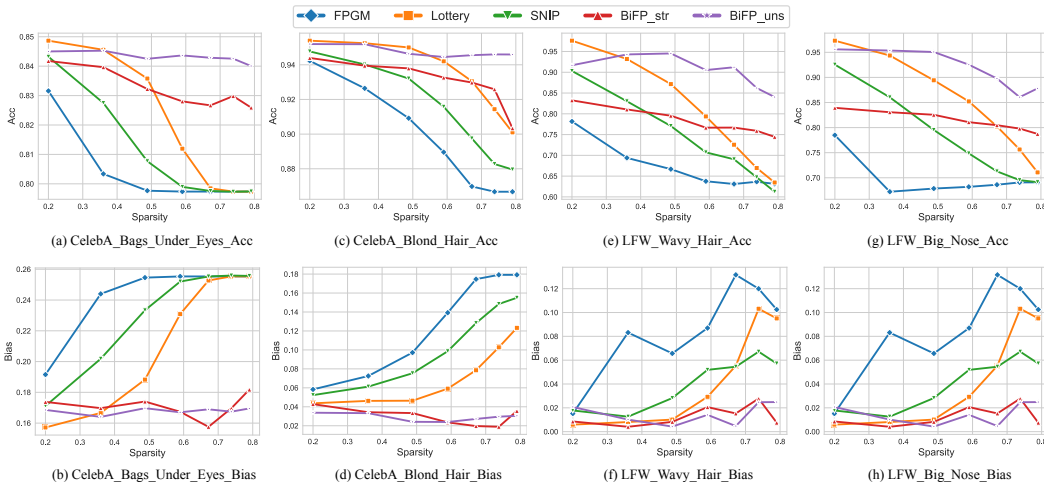


Figure 3: Accuracy and fairness evaluation of pruned **ResNet10** at varying sparsity levels. Subfig. (a, c, e, g) indicates the comparison of accuracy between the proposed method and the baseline methods. Subfig. (b, d, f, h) indicates the comparison of Compressed Model Fairness between the proposed method and the baseline methods. The proposed methods ensure fairness while having comparable even better accuracy.

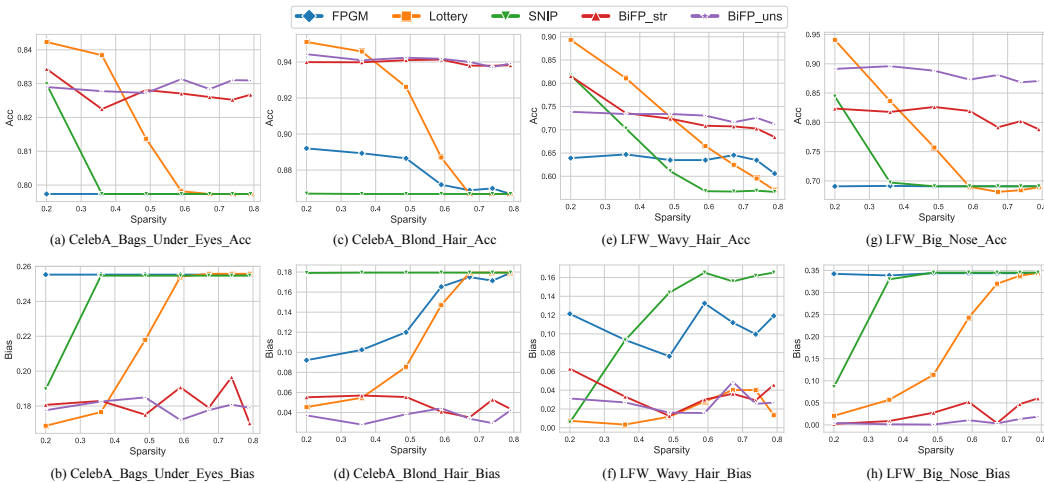


Figure 4: Accuracy and fairness of the pruned **Mobilenetv2** at varying sparsity level. Subfig. (a, c, e, g) indicates the comparison of accuracy between the proposed method and the baseline methods. Subfig. (b, d, f, h) indicates the comparison of Compressed Model Fairness between the proposed method and the baseline methods. The proposed methods ensure fairness while having comparable even better accuracy.

Dataset. We evaluate methods on two real-world image datasets, i.e., CelebA (Liu et al., 2015) and LFW (Huang et al., 2007). We adopt Gender as the sensitive attribute for both two datasets. For the CelebA dataset, we choose Blond Hair and Bags-Under-Eyes as our target attributes. For the LFW dataset, we choose Wavy-Hair and Big-Nose as our target attributes, respectively.

Experiment Settings. We pre-train a ResNet10 (He et al., 2016) network with the training set, prune and fine-tune with the validation set, and then report the accuracy and fairness violations at different sparsity in the test set. We strictly follow the sparsity ratio setting adopted by the Lottery Ticket Hypothesis (LTH) (Frankle & Carbin, 2018) to ensure a fair comparison.

Methods. We deploy the following baseline methods and separate them into two categories: unstructured pruning and structured pruning. Unstructured pruning methods contain the following methods: **Single-shot Network Pruning (SNIP)** calculates the connection sensitivity of edges and prunes those with low sensitivity. The structured pruning method contains **Filter Pruning via Geometric Median (FPGM)** pruning filters with the smallest geometric median. **The Lottery Ticket Hypothesis (Lottery)** states that there is a sparse subnetwork that can achieve similar test accuracy to the original dense network. It is worth pointing out that the proposed **Bi-level Fair Pruning (BiFP)** is capable of different pruning settings. We use **BiFP-str** and **BiFP-uns** to denote structured and unstructured variants, respectively.

5.1 PERFORMANCE AND FAIRNESS

We first perform experiments to show the capability of **BiFP** in mitigating bias during the model pruning. Fig. 3 and Fig. 4 compare the accuracy and bias evaluation of several methods on CelebA and LFW datasets. The first row of accuracy results shows that our methods (**BiFP-str** and **BiFP-uns**) are comparable to conventional pruning techniques with regard to model accuracy. The second row of results shows that conventional pruning techniques, including **FPGM**, **Lottery**, and **SNIP**, indeed introduce more bias with higher sparsity while the proposed methods, **BiFP-str** and **BiFP-uns**, consistently ensure fairness at any sparsity levels.

5.2 TRADE-OFF BETWEEN ACCURACY AND FAIRNESS

In the following experiments, we focus on the structured pruning variant of **BiFP** and drop the subscript for simplicity. We further investigate the trade-off between accuracy and fairness of the proposed method **BiFP** on CelebA with **Blond Hair** as our target attribute at the 70% sparsity level and report the results in Fig. 5. Through the trade-off curves, we conclude that the proposed method achieves a smaller bias compared with the baselines if we consider the same accuracy level (in a horizontal view of the plot). In another aspect, the proposed method achieves higher accuracy if we consider the same fairness level (in a vertical view of the plot). In a nutshell, the proposed method has a better trade-off for accuracy and fairness than the baselines.




Figure 5: Trade-off between accuracy and fairness.

5.3 FAIRNESS AND TRAINING EFFICIENCY

To investigate the fairness capability and training efficiency of the proposed method and baselines, we amend the baselines with a two-stage fair training strategies. In the first stage, we prune a dense model using the conventional pruning techniques to get sparse models. In the second stage, we apply fairness constraint to retrain the parameters of sparse models to achieve the Performance Degradation Fairness. We record the number of training iterations for our proposed solutions and baseline methods and show them in Fig. 6. The red shade indicates the iterations used for **BiFP** to achieve the desired fairness and accuracy. In contrast, the grey shade demonstrates the iterations of the best baseline that has comparable performance to **BiFP**. The gap between the red shade and gray shade indicates that the proposed method is more efficient in achieving fairness. Our proposed method outperforms the best baseline methods while using notably less training cost, specifically savings 94.22% and 94.05% training iterations on the LFW and Celeba datasets.

5.4 ABLATION STUDIES

In the proposed method, we incorporated fairness constraints on the weights and masks of the network pruning simultaneously. To delve deeper into the implications of each component, we conducted an ablation study focusing on the fair weight and fair mask mechanisms separately. In particular, we remove the fair constraints for the mask in Eq. 5 and refer to this variant as **BiFP w/o m**. Similarly, we remove the fair constraints for the weights in Eq. 5 and refer to it as **BiFP w/o w**.




Figure 6: Training iterations used to obtain fair models using a two-stage pruning strategy. The red shade indicates the BiFP iterations while the grey shade indicates iterations of the best baseline.




Figure 7: Ablation Studies on BiFP.




Figure 8: Interpolation of loss surfaces from BiFP to BiFP w/o w (left) & BiFP w/o m (right).

We further remove the fair constraints for both mask and weights, denoted by **BiFP w/o w&m**. The ablation study results are shown in Fig. 7. First, we observe that **BiFP w/o m** and **BiFP w/o w&m** have remarkably lower accuracy and higher bias than the original **BiFP**, indicating masks are necessary for both fairness and fairness. Although **BiFP w/o w** has better accuracy than the original **BiFP** at the low sparsity level, **BiFP** will outperform **BiFP w/o w** in terms of both accuracy and fairness at high sparsity level. The reason is **BiFP w/o w** is capable of preserving accuracy and maintaining fairness merely based on fair weights of the model. If the model is sparse, the weights cannot preserve accuracy and fairness anymore. Additionally, **BiFP w/o w** has relatively high accuracy but high bias as the weights are corrupted.

Loss surface exploration. Furthermore, we explore the implications of **BiFP** by exploring the loss surface of **BiFP** and its counterparts. Fig. 8 shows the loss surface lying between **BiFP** and **BiFP w/o w** or **BiFP w/o m**, respectively. The left plot indicates the loss surface from **BiFP** to **BiFP w/o w** and the right plot indicates the loss surface from **BiFP** to **BiFP w/o m**. We use linear interpolation along the direction between **BiFP** and its counterpart. The right surface plot reveals a high-loss barrier (basin) between **BiFP** and **BiFP w/o m**, which demonstrates the fairness-constrained mask significantly ensures the performance in terms of both accuracy and fairness. It also suggests that only considering fairness-aware weights without fair masks, the optimization can get stuck at local minima that are isolated from better solutions. On the contrary, **BiFP** and **BiFP w/o w** (in the left plot) lie in the same basin, indicating a solution without fair weights is very close to the **BiFP** solution from the optimization perspective. The loss surfaces validate the ablation studies.

6 CONCLUSION

In this paper, we investigate the bias concerns in the neural network pruning task and mathematical formulate fairness notions for this task. Then, we introduce a novel constrained bi-level framework **BiFP** that seamlessly integrates fairness, accuracy, and model sparsity and enables the efficiency of pruning while maintaining both accuracy and fairness. We leverage the bi-level optimization to enforce fairness in both masks and weights simultaneously. In the extensive experiments, we show that the proposed solution is able to mitigate bias and ensure model performance with notably lower training costs.

REFERENCES

- Pranjal Awasthi, Matthäus Kleindessner, and Jamie Morgenstern. Equalized odds postprocessing under imperfect group information. In Silvia Chiappa and Roberto Calandra (eds.), *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, Proceedings of Machine Learning Research, pp. 1770–1780. PMLR, 2020.
- Peter L Bartlett, Michael I Jordan, and Jon D. McAuliffe. Convexity, Classification, and Risk Bounds. *Journal of the American Statistical Association*, (473):138–156, March 2006. ISSN 0162-1459. doi: 10.1198/016214505000000907.
- Cenk Baykal, Lucas Liebenwein, Igor Gilitschenski, Dan Feldman, and Daniela Rus. Data-dependent coresets for compressing neural networks with applications to generalization bounds. *arXiv preprint arXiv:1804.05345*, 2018.
- Toon Calders and Sicco Verwer. Three naive Bayes approaches for discrimination-free classification. (2):277–292, 2010. doi: 10.1007/s10618-010-0190-x.
- Toon Calders, Faisal Kamiran, and Mykola Pechenizkiy. Building classifiers with independency constraints. In Yücel Saygin, Jeffrey Xu Yu, Hillol Kargupta, Wei Wang, Sanjay Ranka, Philip S. Yu, and Xindong Wu (eds.), *ICDM Workshops 2009, IEEE International Conference on Data Mining Workshops, Miami, Florida, USA, 6 December 2009*, pp. 13–18. IEEE Computer Society, 2009. doi: 10.1109/ICDMW.2009.83.
- Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*, pp. 797–806. ACM, 2017. doi: 10.1145/3097983.3098095.
- Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, (2): 295–307, 2015.
- Xuanyi Dong, Junshi Huang, Yi Yang, and Shuicheng Yan. More is less: A more complicated network with less inference complexity. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5840–5848, 2017.
- Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- Yiwen Guo, Anbang Yao, and Yurong Chen. Dynamic network surgery for efficient dnns. *Advances in neural information processing systems*, 2016.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- Victor Petró Bach Hansen and Anders Søgaard. Is the lottery fair? evaluating winning tickets across demographics. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, Findings of ACL, pp. 3214–3224. Association for Computational Linguistics, 2021. doi: 10.18653/v1/2021.findings-acl.284.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 3315–3323, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 770–778. IEEE Computer Society, 2016. doi: 10.1109/CVPR.2016.90.

- Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866*, 2018.
- Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 4340–4349, 2019.
- Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 1389–1397, 2017.
- Torsten Hoefer, Dan Alistarh, Tal Ben-Nun, Nikoli Dryden, and Alexandra Peste. Sparsity in deep learning: Pruning and growth for efficient inference and training in neural networks. *The Journal of Machine Learning Research*, (1):10882–11005, 2021.
- Sara Hooker, Aaron Courville, Gregory Clark, Yann Dauphin, and Andrea Frome. What do compressed deep neural networks forget? *arXiv preprint arXiv:1911.05248*, 2019.
- Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- Eugenia Iofinova, Alexandra Peste, and Dan Alistarh. Bias in pruned vision models: In-depth analysis and countermeasures. *CoRR*, 2023. doi: 10.48550/arXiv.2304.12622.
- Faisal Kamiran and Toon Calders. Classifying without discriminating. In *2009 2nd International Conference on Computer, Control and Communication*, pp. 1–6. IEEE, 2009-02. ISBN 978-1-4244-3313-1. doi: 10.1109/IC4.2009.4909197.
- Faisal Kamiran and Toon Calders. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, (1):1–33, 2012.
- Faisal Kamiran, Toon Calders, and Mykola Pechenizkiy. Discrimination aware decision tree learning. In Geoffrey I. Webb, Bing Liu, Chengqi Zhang, Dimitrios Gunopulos, and Xindong Wu (eds.), *ICDM 2010, the 10th IEEE International Conference on Data Mining, Sydney, Australia, 14-17 December 2010*, pp. 869–874. IEEE Computer Society, 2010. doi: 10.1109/ICDM.2010.50.
- Faisal Kamiran, Asim Karim, and Xiangliang Zhang. Decision theory for discrimination-aware classification. In *Proceedings of the 12nd IEEE International Conference on Data Mining (ICDM 2012)*, pp. 924–929. IEEE, December 2012. ISBN 978-0-7695-4905-7. doi: 10.1109/ICDM.2012.45.
- Yawei Li, Shuhang Gu, Luc Van Gool, and Radu Timofte. Learning filter basis for convolutional neural network compression. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5623–5632, 2019.
- Lucas Liebenwein, Cenk Baykal, Harry Lang, Dan Feldman, and Daniela Rus. Provable filter pruning for efficient neural networks. *arXiv preprint arXiv:1911.07412*, 2019.
- Lucas Liebenwein, Cenk Baykal, Brandon Carter, David Gifford, and Daniela Rus. Lost in pruning: The effects of pruning neural networks beyond test accuracy. *Proceedings of Machine Learning and Systems*, pp. 93–138, 2021.
- Tao Lin, Sebastian U Stich, Luis Barba, Daniil Dmitriev, and Martin Jaggi. Dynamic model pruning with feedback. *arXiv preprint arXiv:2006.07253*, 2020.
- Xiaofeng Lin, Seungbae Kim, and Jungseock Joo. FairGRAPE: Fairness-aware GRADient pruning mEthod for face attribute classification. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XIII*, Lecture Notes in Computer Science, pp. 414–432. Springer, 2022. doi: 10.1007/978-3-031-19778-9_24.

- Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Kwang-Ting Cheng, and Jian Sun. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 3296–3305, 2019.
- Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Charles M. Macal and Arthur P. Hurter. Dependence of bilevel mathematical programs on irrelevant constraints. *Comput. Oper. Res.*, (12):1129–1140, 1997. doi: 10.1016/S0305-0548(97)00025-7.
- David Madras, Elliot Creager, Toniann Pitassi, and Richard S. Zemel. Learning adversarially fair and transferable representations. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, Proceedings of Machine Learning Research, pp. 3381–3390. PMLR, 2018.
- Ayalew Getachew Mersha and Stephan Dempe. Linear bilevel programming with upper level constraints depending on the lower level solution. *Appl. Math. Comput.*, (1):247–254, 2006. doi: 10.1016/j.amc.2005.11.134.
- Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. *arXiv preprint arXiv:1611.06440*, 2016.
- Pavlo Molchanov, Arun Mallya, Stephen Tyree, Iuri Frosio, and Jan Kautz. Importance estimation for neural network pruning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11264–11272, 2019.
- Michela Paganini. Prune responsibly. *arXiv preprint arXiv:2009.09936*, 2020.
- Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What’s hidden in a randomly weighted neural network? In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 11890–11899. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPR42600.2020.01191. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Ramanujan_Whats_Hidden_in_a_Randomly_Weighted_Neural_Network_CVPR_2020_paper.html.
- Ao Ren, Tianyun Zhang, Shaokai Ye, Jiayu Li, Wenyao Xu, Xuehai Qian, Xue Lin, and Yanzhi Wang. Admm-nn: An algorithm-hardware co-design framework of dnns using alternating direction methods of multipliers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 925–938, 2019.
- Tiago Salvador, Stephanie Cairns, Vikram Voleti, Noah Marshall, and Adam M. Oberman. FairCal: Fairness calibration for face verification. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- Vikash Sehwal, Shiqi Wang, Prateek Mittal, and Suman Jana. Hydra: Pruning adversarially robust neural networks. *Advances in Neural Information Processing Systems*, pp. 19655–19666, 2020.
- Pengwei Tang, Wei Yao, Zhicong Li, and Yong Liu. Fair scratch tickets: Finding fair sparse networks without weight training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 24406–24416, June 2023.
- Cuong Tran, Ferdinando Fioretto, Jung-Eun Kim, and Rakshit Naidu. Pruning has a disparate impact on model accuracy. In *NeurIPS*, 2022.
- Fu-En Wang, Chien-Yi Wang, Min Sun, and Shang-Hong Lai. MixFairFace: Towards ultimate fairness via MixFair adapter in face recognition. In Brian Williams, Yiling Chen, and Jennifer Neville (eds.), *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pp. 14531–14538. AAAI Press, 2023.

- Yongkai Wu, Lu Zhang, and Xintao Wu. On convexity and bounds of fairness-aware classification. In Ling Liu, Ryen W. White, Amin Mantrach, Fabrizio Silvestri, Julian J. McAuley, Ricardo Baeza-Yates, and Leila Zia (eds.), *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pp. 3356–3362. ACM, 2019. doi: 10.1145/3308558.3313723.
- Jianbo Ye, Xin Lu, Zhe Lin, and James Z Wang. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *arXiv preprint arXiv:1802.00124*, 2018.
- Mao Ye, Chengyue Gong, Lizhen Nie, Denny Zhou, Adam Klivans, and Qiang Liu. Good sub-networks provably exist: Pruning via greedy forward selection. In *International Conference on Machine Learning*, pp. 10820–10830. PMLR, 2020.
- Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and Larry S Davis. Nisp: Pruning networks using neuron importance score propagation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9194–9203, 2018.
- Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P. Gummadi. Fairness constraints: Mechanisms for fair classification. In Aarti Singh and Xiaojin (Jerry) Zhu (eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, Proceedings of Machine Learning Research, pp. 962–970. PMLR, 2017.
- Tianyun Zhang, Shaokai Ye, Kaiqi Zhang, Jian Tang, Wujie Wen, Makan Fardad, and Yanzhi Wang. A systematic dnn weight pruning framework using alternating direction method of multipliers. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 184–199, 2018.
- Yihua Zhang, Yuguang Yao, Parikshit Ram, Pu Zhao, Tianlong Chen, Mingyi Hong, Yanzhi Wang, and Sijia Liu. Advancing model pruning via bi-level optimization. *Advances in Neural Information Processing Systems*, pp. 18309–18326, 2022.
- Junjie Zhu, Lin Gu, Xiaoxiao Wu, Zheng Li, Tatsuya Harada, and Yingying Zhu. People taking photos that faces never share: Privacy protection and fairness enhancement from camera to user. In Brian Williams, Yiling Chen, and Jennifer Neville (eds.), *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, pp. 14646–14654. AAAI Press, 2023.
- Indre Zliobaite. Measuring discrimination in algorithmic decision making. (4):1060–1089, 2017. doi: 10.1007/s10618-017-0506-1.

A APPENDIX

You may include other additional sections here.