



NEUFAIR: Neural Network Fairness Repair with Dropout

Vishnu Asutosh Dasu

Pennsylvania State University

State College, USA

vdasu@psu.edu

Saeid Tizpaz-Niari

University of Texas at El Paso

El Paso, USA

saeid@utep.edu

Ashish Kumar

Pennsylvania State University

State College, USA

azk640@psu.edu

Gang Tan

Pennsylvania State University

State College, USA

gtan@psu.edu

Abstract

This paper investigates neuron dropout as a post-processing bias mitigation method for deep neural networks (DNNs). Neural-driven software solutions are increasingly applied in socially critical domains with significant fairness implications. While DNNs are exceptional at learning statistical patterns from data, they may encode and amplify historical biases. Existing bias mitigation algorithms often require modifying the input dataset or the learning algorithms. We posit that prevalent dropout methods may be an effective and less intrusive approach to improve fairness of pre-trained DNNs during inference. However, finding the ideal set of neurons to drop is a combinatorial problem.

We propose NEUFAIR, a family of post-processing randomized algorithms that mitigate unfairness in pre-trained DNNs via dropouts during inference. Our randomized search is guided by an objective to minimize discrimination while maintaining the model's utility. We show that NEUFAIR is efficient and effective in improving fairness (up to 69%) with minimal or no model performance degradation. We provide intuitive explanations of these phenomena and carefully examine the influence of various hyperparameters of NEUFAIR on the results. Finally, we empirically and conceptually compare NEUFAIR to different state-of-the-art bias mitigators.

CCS Concepts

- Software and its engineering → Search-based software engineering;
- Computing methodologies → Machine learning.

Keywords

Machine Learning, Bias Mitigation, AI Ethics

ACM Reference Format:

Vishnu Asutosh Dasu, Ashish Kumar, Saeid Tizpaz-Niari, and Gang Tan. 2024. NEUFAIR: Neural Network Fairness Repair with Dropout. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA '24)*, September 16–20, 2024, Vienna, Austria. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3650212.3680380>

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

ISSTA '24, September 16–20, 2024, Vienna, Austria

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0612-7/24/09

<https://doi.org/10.1145/3650212.3680380>

1 Introduction

Artificial intelligence (AI), increasingly deployed with deep neural network (DNN) components, has become an integral part of modern software solutions that assist in socio-economic and legal-critical decision-making processes such as releasing patients [31], identifying loan defaults [23], and detecting tax evasion [43].

Despite many advances made possible by AI, some challenges require understanding the dimensions and implications of deploying AI-driven software solutions. One such concern about the trustworthiness of AI is discrimination. Unfortunately, there are plenty of fairness defects in real systems. Parole decision-making software was found to harm black and Hispanic defendants by falsely predicting a higher risk of recidivism than for non-Hispanic white defendants [28]; Amazon's hiring algorithm disproportionately rejected more female applicants than male applicants [27]; and data-driven auditing algorithm selected black taxpayers with earned income tax credit claims (EITC) at much higher rates than other racial groups for an audit [43]. As evidenced by these examples, resulting software may particularly disadvantage minorities and protected groups and be found non-compliant with laws such as the US Civil Rights Act [6]. Hence, helping programmers and users to mitigate unfairness in social-critical data-driven software systems is crucial to ensure inclusion in our modern, increasingly digital society.

The software engineering (SE) community has spent significant efforts to address discrimination in the automated data-driven software solutions [2, 10, 19, 47, 56]. Fairness has been treated as a critical meta-property that requires analysis beyond functional correctness and measurements beyond prediction accuracy [7]. Thus, the community presents various testing [2, 19, 57], debugging [22, 36, 50], and mitigation [5, 53] techniques to address fairness defects in data-driven software.

Broadly, fairness mitigation can be applied in the pre-processing (e.g., increasing the representations of an under-represented group by generating more data samples for them), in-processing (e.g., changing the loss function to include fairness constraints during training), or post-processing (e.g., changing the logic of a pre-trained model) stage. However, when the decision logic of AI is encoded via DNNs, it becomes challenging to mitigate unfairness due to its black-box uninterpretable nature.

We posit that a subset of neurons in a neural network disproportionately contributes to unfairness. Removing these neurons during inference as a post-processing operation on a trained DNN can improve fairness. Therefore, prevalent techniques such as dropout

methods [40, 41]—a process of randomly dropping neurons during training—might be effective in mitigating unfairness in *pre-trained* DNN models. While the dropout strategy has been significantly used to prevent over-fitting during the training of DNNs, to the best of our knowledge, this is the first work to systematically leverage dropout methods as a post-processing bias mitigation method for improving the fairness of *pre-trained* DNN models.

However, finding the optimal subset of dropout neurons is an intractable combinatorial problem that requires an exhaustive search over all possible combinations of neuron dropouts. To overcome the computational challenges, we explore the class of randomized algorithms with Markov chain Monte Carlo (MCMC) strategies to efficiently explore the search space with statistical guarantees. In doing so, we pose the following research questions:

RQ1. How successful are randomized algorithms in repairing the unfairness of DNNs via dropouts?

RQ2. Are there dropout strategies that improve model fairness and utility together?

RQ3. What are the design considerations of search algorithms for efficient and effective dropouts?

RQ4. How do dropout strategies compare to the state-of-the-art post-processing (bias) mitigators?

To answer these research questions, we present NEUFAIR (Neural Network Fairness Repair): a set of randomized search algorithms to improve the fairness of DNNs via inference-time neuron dropouts. We design and implement simulated annealing (SA) and random walk (RW) strategies that efficiently explore the state-space of neuron dropouts where we encode the frontiers of fairness and utility in a cost function.

We evaluate NEUFAIR over 7 deep neural network benchmarks trained over 5 socially critical applications with significant fairness implications. We found that NEUFAIR can improve fairness (up to 69%) with minimal utility degradation in most cases. We also report a pathological case and reasons behind a failure of NEUFAIR. We also observe that NEUFAIR can simultaneously improve fairness and utility and provide intuitive explanations of such phenomena. Furthermore, we examine different hyperparameter configuration options of randomized algorithms. While some hyperparameters always influence fairness with positive or negative impacts, we detect a hyperparameter that defines a trade-off between explorations and exploitation that should be tuned as a constant variable in the cost function for each benchmark. Finally, we show the effectiveness of the SA algorithm, compared to the RW and the state-of-the-art post-processing (bias) mitigator [36].

In summary, the key contributions of this paper are:

- (1) **Inference-Time dropout for fairness.** To the best of our knowledge, we present the first dropout method of bias mitigation over pre-trained deep neural networks,
- (2) **Randomized algorithms for fairness.** We create a well-defined framework to formulate the combinatorial inference-time dropout problem in DNNs using randomized algorithms,
- (3) **Experimental evaluations.** We implement the randomized algorithms in NEUFAIR and evaluate their effectiveness and efficiency vis-a-vis the state-of-the-art techniques.

2 Background

In this section, we provide a background on the various model utility and fairness metrics.

2.1 Notions of Model Utility

Given a binary classifier h , a set of features X , and predictions $h(X) = \{TP, TN, FP, FN\}$, we can define the following notions of model utility. In $h(X)$, TP, TN, FP, and FN denote the set of True Positives, True Negatives, False Positives, and False Negatives. We drop the cardinality operator $|\cdot|$ in the following definitions for brevity.

$$\text{Accuracy}_h = \frac{TP + TN}{TP + TN + FP + FN}$$

$$F1_h = \frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}} = \frac{2 * TP}{2 * TP + FP + FN}$$

Accuracy can be used to gauge the overall performance of a classifier. However, accuracy is a poor metric for imbalanced datasets commonly used in fairness evaluations as the number of negative samples far outweighs the positive samples. For example, predicting all samples as negative (0 true positives) in the *Bank* [14] dataset yields an accuracy of 88% but the F1 score would be 0 or undefined. Since F1 is defined as the harmonic mean of precision and recall, rather than the arithmetic mean, it penalizes performance significantly if either precision or recall is low.

2.2 Notions of Model Fairness

Consider a machine learning classifier h , a set of features X , sensitive features $A \subset X$, and a set of labels Y . We can then define the following notions of fairness.

DEFINITION 2.1 (DEMOGRAPHIC PARITY [1]). *The classifier h satisfies Demographic Parity under a distribution over (X, A, Y) if its prediction $h(X)$ is statistically independent of the sensitive feature A i.e. $\mathbb{P}[h(X) = \hat{y}|A = a] = \mathbb{P}[h(X) = \hat{y}]$ for all a, \hat{y} . For binary classification with $\hat{y} = \{0, 1\}$, this is equivalent to $\mathbb{E}[h(X)|A = a] = \mathbb{E}[h(X)]$ for all a .*

DEFINITION 2.2 (EQUALIZED ODDS [1, 25]). *The classifier h satisfies Equalized Odds under a distribution over (X, A, Y) if its prediction $h(X)$ is conditionally independent of the sensitive feature A given the label Y i.e. $\mathbb{P}[h(X) = \hat{y}|A = a, Y = y] = \mathbb{P}[h(X) = \hat{y}|Y = y]$ for all a, y , and \hat{y} . For binary classification with $\hat{y} = \{0, 1\}$, this is equivalent to $\mathbb{E}[h(X)|A = a, Y = y] = \mathbb{E}[h(X)|Y = y]$ for all a, y .*

Equal Opportunity is a relaxed variant of Equalized Odds with $Y = 1$ [25]. Equalized Odds require the true positive and false positive rates to be equal across all sensitive groups. However, Equal Opportunity only requires the true positive rate to be equal across all sensitive groups.

Demographic Parity is the weakest notion of fairness, and Equalized Odds is the strongest. In our work, we use Equalized Odds as the fairness criterion. The disparity or unfairness for Equalized Odds is the Equalized Odds Difference (EOD), defined as the maximum absolute difference between the true and false positive rates across the sensitive groups. Mathematically, this is represented as

$$EOD = \max \left(\begin{array}{l} |\mathbb{P}[h(X) = 1 | A = 0, Y = 1] - \mathbb{P}[h(X) = 1 | A = 1, Y = 1]|, \\ |\mathbb{P}[h(X) = 1 | A = 0, Y = 0] - \mathbb{P}[h(X) = 1 | A = 1, Y = 0]| \end{array} \right)$$

with sensitive features $A = \{0, 1\}$.

3 Problem Statement

We consider pre-trained deep neural network (DNN) classifiers with the set of input variables X partitioned into a protected set of variables (such as race, sex, and age) and non-protected variables (such as profession, income, and education). We further assume the output is a binary classifier that gives either favorable or unfavorable outcomes.

3.1 Syntax and Semantics of DNN

A deep neural network (DNN) encodes a function $\mathcal{D} : X \rightarrow [0, 1]^2$ where X consists of the set of protected attributes $X_1 \times X_2 \times \dots \times X_m$ and non-protected attributes $X_{m+1} \times \dots \times X_{m+r}$. The DNN model is parameterized by the input dimension $m+r$, the output dimension 2, the depth of hidden layers $n+1$, and the weights of its hidden layers W_0, W_1, \dots, W_n . We describe the hidden layers with $\mathcal{M} \leftarrow [L_I, L_0, \dots, L_n, L_O]$, where L_I and L_O are the input and output layers, respectively, and $L_i, \forall i \in [0, n]$, are the hidden layers. We assume that there exists a subset of neurons $\mathcal{N} \in L_i, \forall i \in [0, n]$ in the hidden layers that disparately contribute to unfairness.

Let L_i be the output of layer i that implements an affine mapping from the output of previous layer L_{i-1} and its weights W_{i-1} for $1 \leq i \leq n$ followed by a fixed non-linear activation unit (e.g., ReLU defined as $L_{i-1} \mapsto \max\{W_{i-1} \cdot L_{i-1}, 0\}$) for $1 \leq i \leq n$. Let L_i^j be the output of neuron j at layer i that is $L_i^j(x) = \text{ReLU}\left(\sum_{j=1}^{|L_i|} w_i L_{i-1}^j\right)$. The output is the likelihood of favorable and unfavorable outcomes. The predicted label is the index of the maximum likelihood, $\mathcal{D}(x) = \max_i L_O(x)(i)$.

3.2 Inference Time Dropout for Fairness

Dropout [41] is a technique proposed to improve the performance of DNNs by preventing overfitting. Dropout sets all w_i to 0 for a random set of neurons in the hidden layers with some probability during training. Once training is complete, dropout is not used, and all the neurons in the DNN are utilized to make predictions. While dropout has been traditionally used to prevent overfitting during training, we hypothesize that dropping neurons of the DNN during inference after training can significantly improve fairness with a minimal impact on performance. However, unlike traditional dropout, where a set of neurons are randomly dropped during training, we aim to identify a subset of neurons at the Pareto optimal curve of fairness-performance during inference.

We consider a binary vector as the neuron state with $s = \{0, 1\}^N$, where $N = \sum_{i=0}^n |L_i|$ and s_i indicates whether the neuron i is dropped or not and n is the number of layers. A pre-trained DNN model \mathcal{D} does not include any dropouts; hence, all the indicators are 0.

DEFINITION 3.1 (DESIRABLE DROPOUT OF FAIRNESS VS. UTILITY). *Given a DNN model \mathcal{D} trained over a dataset X ; the search problem is to infer a repaired DNN model \mathcal{D}' by dropping a subset of neurons*

I in the binary neuron state, i.e., $s_i = 1$ for any $i \in I$, such that (1) the model bias (e.g., EOD) is maximally reduced, (2) the model performance (e.g., F1-score) is minimally degraded, and (3) the model structure in terms of the numbers of inputs, outputs, and hidden layers remains the same as compared to the original pre-trained model \mathcal{D} .

We define desirable states as those states that have a good fairness-performance tradeoff. A brute-force search to find desirable neurons is exponential in the size of DNN as we have $2^{|N|}$ possible subsets. The running time becomes prohibitively expensive, even for small DNNs. To find the desirable subset of neurons, we explore different types of randomized algorithms to improve fairness via model inference dropout.

4 Approach

Our approach comprises two randomized search algorithms, namely Simulated Annealing (SA) and Random Walk (RW).

4.1 Simulated Annealing Search

We formulate the problem of finding the desirable subset of neurons as a discrete optimization problem and solve it using Simulated Annealing. Simulated Annealing (SA) [4] is a probabilistic algorithm that finds the global minima (w.r.t some cost function) in a large search space with high probability. Algorithm 1 presents a generic template to apply SA to a search space optimization problem. Figure 1 overviews the steps in our bias mitigation approach with the SA search. We now define the core concepts used in our SA algorithm.

Algorithm 1: Generic Simulated Annealing (SA) Procedure

```

Input: Search space  $\mathcal{S}$ , neighborhood relation  $\Gamma$ , cost
       function  $c$ , initial temperature  $T_0$ , initial state  $s_0$ 
Output: Best solution found.

1  $T, s, s^*, m \leftarrow T_0, s_0, s_0, 0$ 
2 while 'stopping criterion' is not met do
3    $s' \leftarrow \text{Generate}(s)$ 
4   Compute  $\Delta E = \text{cost}(s') - \text{cost}(s)$ 
5   if  $\Delta E \leq 0$  then
6     Accept transition i.e.  $s \leftarrow s'$ 
7     Update best solution found i.e.  $s^* \leftarrow s'$ 
8   else
9     if  $e^{-\Delta E/T_m} \geq \text{UniformSample}(0, 1)$  then
10       Accept transition i.e.  $s \leftarrow s'$ 
11    $T_{m+1} \leftarrow \text{Update}(T_m)$ 
12    $m \leftarrow m + 1$ 
13 return  $s^*$ 

```

Search Space. For a DNN $\mathcal{M} \leftarrow [L_I, L_0, \dots, L_n, L_O]$ we define a state of our search space \mathcal{S} as a binary sequence $s \in \{0, 1\}^N$, where $N = \sum_{i=0}^n |L_i|^1$. The i th element of s , denoted s_i is 1 if neuron i is

¹For memory reasons, we store a state as a decimal number rather than a N -bit binary sequence, where the binary sequence s is mapped to the unique decimal number which has binary expansion 1s.

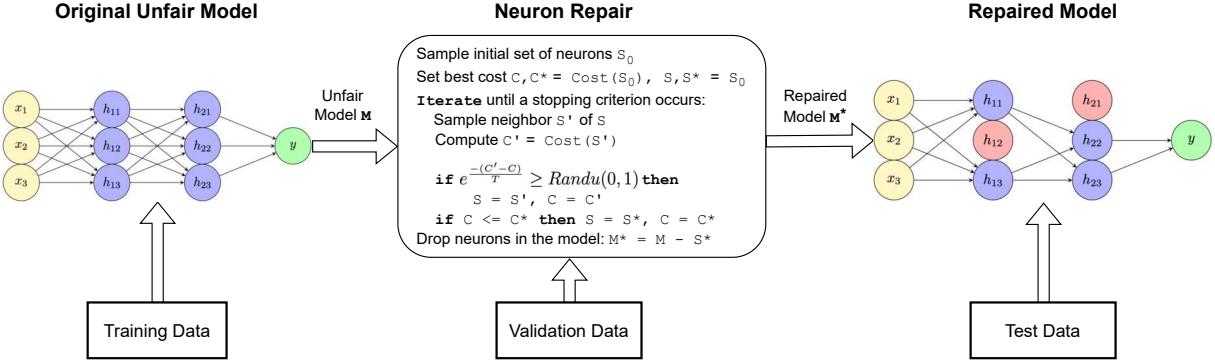


Figure 1: Overview of the NEUFAIR: a post-processing bias mitigation approach.

dropped and 0 otherwise². For example, consider a DNN $\mathcal{M} \leftarrow [L_I, L_0, L_1, L_O]$, where $|L_0| = 3$ and $|L_1| = 3$. Then, the state s given by:

$$s = (0, \underbrace{1, 0, 0}_L, \underbrace{0, 0, 1}_L)$$

drops the second neuron in the first layer and the third neuron in the second layer.

Instead of allowing our search space \mathcal{S} to drop every possible subset of neurons possible (which would ensure \mathcal{S} has size 2^N), we restrict the size of \mathcal{S} by fixing an upper and lower bound on the number of neurons that can be dropped from the DNN. Let n_l and n_u ($n_l \leq n_u$) denote the minimum and maximum number of neurons allowed by our DNN, respectively. We can then formally define \mathcal{S} as:

$$\mathcal{S} := \{s \in \{0, 1\}^N \mid n_l \leq HW(s) \leq n_u\} \quad (1)$$

where $HW(s)$ denotes the hamming weight of s . Restricting our search space with a conservative estimate of the lower and upper bound would have a minimal impact on our ability to find the desirable subset, as dropping too many neurons will reduce the model utility to less than acceptable levels, whereas dropping too few neurons will improve fairness by only a marginal amount. However, asymptotically, a conservative estimate of the lower and upper bound still generates prohibitively large search spaces. With the bounds n_l and n_u , the cardinality of the search space $|\mathcal{S}| = \sum_{i=n_l}^{n_u} \binom{N}{i} = \Omega(\frac{2^N}{\sqrt{N}})$ if $n_l < \frac{N}{2} < n_u$, which rules out brute-force as a viable option.

Neighborhood Relation and ‘Generate’ subroutine. The neighborhood of any state $s \in \mathcal{S}$, denoted $\Gamma(s)$, is defined as the set of all states that are at a hamming distance of 1. Mathematically, this is defined as

$$\Gamma(s) := \{s' \in \mathcal{S} \mid HD(s, s') = 1\}$$

where $HD(s, s')$ denotes the hamming distance between s and s' . With our definition of the search space and neighborhood of a state, the entire search space graph can be viewed as a subset of the

²The neurons are numbered according to some total order; the choice of total order is insignificant as long as we fix the mapping of a neuron to its position in the binary sequence.

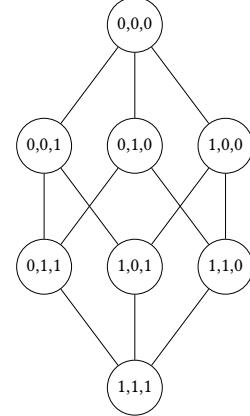


Figure 2: 3-dimensional Boolean hypercube that demonstrates the exponential search space of a neural network with 1 hidden layer having 3 neurons.

N -dimensional hypercube, where the vertices of the hypercube represent the states of our search space, and the edges of the hypercube represent the neighborhood relation. The *Generate* subroutine on input $s \in \mathcal{S}$ uniformly samples a neighbor s' from $\Gamma(s)$ and returns s' . This is equivalent to uniformly sampling an index position i from $[1, n]$ and subsequently uniformly flipping bit s_i in the binary sequence s to get s' . We make the following observation on our underlying search space graph³:

LEMMA 4.1. *If $n_l < n_u$, then our search space graph is connected and has a diameter less than N . Moreover, the distance between any two states $s, s' \in \mathcal{S}$ is given by $HD(s, s')$.*

For example, Figure 2 shows the search space graph for a neural network with 1 hidden layer with 3 neurons, along with the neighbors and state transitions.

Cost Function. Our primary goal is to find a state $s \in \mathcal{S}$ which minimizes its unfairness score EOD_s . However, we do not wish to

³The search space graph is a graph with \mathcal{S} as the set of vertices and edges defined according to the neighborhood relation Γ i.e. (s, s') is an edge iff $s' \in \Gamma(s)$. As our neighborhood relation is symmetric, our search space graph is undirected.

consider those states with a significant loss in model performance (measured using the F1 score of the model) compared to the original model; thus, we try to find an acceptable balance between the improvement in fairness and the loss in model performance. To realize this balance, we penalize our cost function with an additional term to artificially increase the cost if the state has a lower F1 score. We now formally define our cost function, $\text{cost}(\cdot)$ as follows:

$$\text{cost}(s) := EOD_s + p \cdot EOD_{s_0} \cdot \mathbb{1}(F1_s < tF1_{s_0}) \quad (2)$$

In the above equation, s_0 is the original state (the DNN with no dropout), $p \in \mathbb{R}_{\geq 0}$ is called the penalty multiplier, $t \in (0, 1)$ is called the threshold multiplier, EOD_s and EOD_{s_0} are the unfairness scores of states s and s_0 respectively, $F1_s$ and $F1_{s_0}$ are the F1 scores of the s and s_0 respectively, and $\mathbb{1}(\cdot)$ denotes the indicator function. The threshold multiplier t determines the percentage loss in F1 score we will tolerate to improve fairness. The penalty multiplier p discourages states with an F1 score less than the threshold by penalizing them with a multiple of the unfairness of the initial state. Our cost function formulation allows us to find the state with the minimum unfairness while maintaining a significantly higher F1.

Initial Temperature, Cooling Schedule, and ‘Update’ subroutine. The temperature T of the SA procedure determines the probability with which we accept a positive transition (a transition where the cost is increased); the higher the temperature, the more likely the algorithm accepts such a transition. The cooling schedule refers to the function used to update the temperature after each iteration. We adopt the logarithmic cooling schedule [24] which has proven convergence guarantees [35]. According to the logarithmic cooling schedule, the temperature $T(\cdot)$ at iteration m is defined as

$$T_m = \frac{T_0}{\log(2 + m)}, \forall m \in \mathbb{Z}_{\geq 0} \quad (3)$$

The *Update* subroutine updates the temperature using the above formula. Using the convergence results for SA with a logarithmic cooling schedule [35], we get the following result for our SA runs:

LEMMA 4.2. *Let $p > 1$. If we set $T_0 \geq (1 + p \cdot EOD_{s_0})(n_u - n_l)$, then the probability that SA finds a global minima within $k > \lceil \frac{n_u - n_l}{2} \rceil$ iterations is greater than $1 - \frac{A}{(\frac{k}{n_u - n_l})^c}$ where*

$$c := \min \left(\frac{1}{(n_u - n_l)N^{\lceil \frac{n_u - n_l}{2} \rceil}}, \frac{4}{T_0 S^2} \right)$$

where S is the size of the validation dataset and $A > 0$ is a constant.

While the above lemma provides a worst-case bound on the number of iterations required to find an optimal solution, it is practically infeasible to achieve the large bound. Ben-Ameur [3] proposes an algorithm for computing T_0 and shows that SA performs well experimentally with fewer iterations than estimated by Mitra et al. [35]. We use the temperature initialization algorithm from [3] to estimate the initial temperature T_0 . At a high level, this algorithm back-computes an initial value for the temperature for which the expected value of acceptance probabilities for positive transitions from a random initial distribution is greater than some predefined threshold.

Stopping Criterion. We use time as a stopping criterion for our SA runs i.e., we run the SA search for time less than some hyperparameter *time_limit*.

Algorithm 2: NEUFAIR to mitigate unfairness in trained neural networks

Input: Unfair neural network \mathcal{M} , Penalty multiplier p , Threshold multiplier t , Minimum and maximum number of neurons to drop $[n_l, n_u]$, Algorithm Type *alg_type*, Time Limit *time_limit*

Output: Repaired neural network \mathcal{M}_\star , Desirable state s_\star , Best cost cost_\star

```

1  $s \leftarrow \text{random\_state}(\mathcal{M}, n_l, n_u)$ 
2  $s_\star, \text{start\_time} \leftarrow \phi, \text{curr\_time}()$ 
3  $\text{cost} \leftarrow \text{compute\_cost}(\mathcal{M}, s, p, t)$ 
4  $\text{cost}_\star \leftarrow \text{compute\_cost}(\mathcal{M}, s_\star, p, t)$ 
5  $T_0 \leftarrow \text{estimate\_temperature}(\mathcal{M}, s)$ 
6 while  $\text{curr\_time}() - \text{start\_time} \leq \text{time\_limit}$  do
7    $T \leftarrow \text{update\_temperature}(T_0, \text{curr\_time}())$ 
8    $s_i \leftarrow \text{generate\_state}(s, n_l, n_u)$ 
9    $\text{cost}_i \leftarrow \text{compute\_cost}(\mathcal{M}, s_i, p, t)$ 
10   $\Delta E \leftarrow \text{cost}_i - \text{cost}$ 
11  if  $\Delta E \leq 0$  then
12     $\text{cost} \leftarrow \text{cost}_i$ 
13     $s \leftarrow s_i$ 
14  else if (alg_type == RW)  $\vee$  (alg_type ==
15    SA  $\wedge e^{-\Delta E/T} \geq \text{Uniform}(0,1)$ ) then
16     $\text{cost} \leftarrow \text{cost}_i$ 
17     $s \leftarrow s_i$ 
18  if  $\text{cost} \leq \text{cost}_\star$  then
19     $\text{cost}_\star \leftarrow \text{cost}_i$ 
     $s_\star \leftarrow s_i$ 
20  $\mathcal{M}_\star \leftarrow \mathcal{M} \setminus s_\star$ 
21 return  $\mathcal{M}_\star, s_\star, \text{cost}_\star$ 

```

4.2 Random Walk Search

The Random Walk (RW) strategy samples a random state $s \in \mathcal{S}$, and recursively samples states from the neighborhood $\Gamma(\cdot)$ to explore the search space. The RW strategy then records the best state discovered throughout the walk. We use the same cost function as SA in RW to determine the desirable state. A key difference between RW and SA search is that in RW, we always transition to a new state, regardless of cost. In contrast, in SA, we always transition to a new state with a lower cost and transition with some probability if it has a higher cost. This is highlighted in Line 14 in Algorithm 2. A RW can also be considered an SA run with infinite temperature, i.e., the transition probability is always 1.

5 Experiments

We pose the following research questions:

RQ1 How successful are randomized algorithms in repairing the unfairness of DNNs via dropouts?

Table 1: Architectures of the DNNs used for the datasets.

Dataset	Model Architecture [$L_1, L_0, \dots, L_n, L_O$]
Adult Census Income	[34, 64, 128, 64, 1]
Compas Software	[12, 32, 32, 1]
Bank Marketing	[32, 32, 32, 1]
Default Credit	[30, 16, 16, 16, 1]
Medical Expenditure (MEPS16)	[138, 128, 128, 128, 1]

RQ2 Are there dropout strategies that improve fairness and utility together?

RQ3 What are the design considerations of search algorithms for efficient and effective fairness repair of DNNs via dropout?

RQ4 How do dropout strategies compare to the state-of-the-art post-processing (bias) mitigators?

5.1 Datasets and Models

We evaluate NEUFAIR with five different datasets from fairness literature. For two of the datasets, we consider two different protected groups, which effectively results in a total of 7 benchmarks. Table 2 presents an overview of the datasets and protected groups. The *Adult* Census Income [15], *Bank* Marketing [14], *Compas* Software [39], *Default* Credit [16], and Medical Expenditure (*MEPS16*) [18] are binary classification tasks to predict whether an individual has income over 50K, is likely to subscribe, has a low reoffending risk, is likely to default on the credit card payment, and is likely to utilize medical benefits, respectively.

We used DNNs trained over the dataset benchmarks as the machine learning model. Table 1 highlights the architectures of the DNN models used for each dataset. We use ReLU as the activation function after the linear layers. For the *Bank*, *Default*, and *Compas* datasets we use a dropout of 0.2 during training. For the *Adult* and *MEPS16* datasets, we set the dropout to 0.1. For the *Compas* dataset, we use Adam as the optimizer with a learning rate of 0.001. For the other datasets, we use SGD with a learning rate of 0.01. For data preprocessing, we utilize standard techniques such as one-hot encoding for categorical features followed by min-max or standard scaling for numerical features. For the *Compas* dataset, we use a version of the dataset used in [46] that has 12 features after feature selection. The modified *Compas* dataset is available online⁴.

5.2 Technical Details

We run all our experiments on a desktop running Ubuntu 22.04.3 LTS with an Intel(R) Core(TM) i7-7700 CPU @ 3.60 GHz processor, 32 GB RAM, and a 1 TB HDD. The neural networks and repair algorithms were implemented using `python3.8`, `torch==2.0.1`, `numpy==1.24.3`, and `scikit-learn==1.3.1`.

5.3 Experimental Setup

For each dataset, we evaluate the performance of the NEUFAIR algorithms using 10 different random seeds. The seeds are set for the `torch`, `numpy`, and `scikit-learn` libraries before DNN training.

⁴<https://github.com/Tizpaz/Parfait-ML/blob/main/subjects/datasets/compas>

Specifically, the seed determines the randomness of the training, validation, and test splits and the randomness of model training, such as sampling batches and initializing the weights of the DNN. The NEUFAIR algorithms themselves are not seeded. Each SA and RW run is unique.

During training, the training dataset was used for gradient descent, and the validation dataset was used for hyperparameter tuning. We use the model from the epoch with the highest validation F1 score. The validation dataset is also used by NEUFAIR to identify the desirable set of neurons to drop. We evaluate the fairness of the original and fixed (repaired) models using the test dataset.

Unless otherwise specified, all runs reported in the paper have a time-out limit of 1 hour with the following cost function:

$$C(s) = EOD_s + 3.0 \cdot EOD_{s_0} \cdot \mathbf{1}(F1_s < (0.98 \cdot F1_{s_0})).$$

The F1 threshold is 98% of the validation F1 score of the unfair DNN, i.e., we tolerate a 2% degradation in the F1 score to improve fairness. A penalty of 3 \times the baseline unfairness is added to the fairness of the current state if its F1 score is less than the threshold. For all datasets, we set the minimum number of neurons to $n_l = 2$ and vary the maximum number of neurons n_u between 20% – 40% of the number of hidden layer neurons. The n_u values are 50, 24, 24, 20, and 135 for *Adult*, *Compas*, *Bank*, *Default*, and *MEPS16* respectively. For all datasets, we use a train/validation/test split of 60%/20%/20%. We set the threshold for the acceptance probabilities in the temperature initialization algorithm [3] to 0.75.

5.3.1 Effectiveness of randomized algorithms on mitigating unfairness using dropout (RQ1). The first research question is answered by analyzing the data in Figure 3, Figure 4, Table 3, and Table 4.

Figure 3 and Figure 4 highlight the cost of the best state found for the SA and RW over time, respectively, for 10 different seeds per data subject using the *logarithmic scale*. In all cases, we see a positive trend where the cost improves over time, and the search identifies better subsets of neurons to drop. We note the seemingly flat plots correspond to those cases where a good state was found in the beginning. The exception is the black line in the sub-plot for the *Default* dataset in both figures as it corresponds to a failed run; i.e., the best state found has an F1 score lesser than 0.98 \times the F1 score of the initial state. Out of 70 runs (10 per benchmark), the SA and RW found a desirable state with 0.98 \times the initial F1 score in 69 cases. In addition, we re-run the failed *Default* experiment with an F1 threshold multiplier of 0.96 and identify a desirable state with a test set fairness of 6.26%, an improvement from 13.2%. The progress graphs of the RW strategy are “steep” initially as the initial states are unfair and are heavily penalized by the penalty multiplier as RW generates intermediate states with an F1 score less than the threshold. However, due to the cost function and probabilistic transitions, SA only explores worse states with a low probability. Thus, the progress graphs are not as steep initially.

Table 3 and Table 4 highlight the mean and 95% confidence intervals for the F1 score, accuracy, and EOD for all data subjects aggregated across 10 seeds. We observe that in all cases, both NEUFAIR algorithms improve the EOD. However, SA outperformed RW in all cases for the validation set while ensuring that the F1 score reduces by at most 2%. We can conclude that SA is the more

Table 2: Datasets used in our experiments.

Dataset	Instances	Features	Protected Groups		Outcome Label	
			Group1	Group2	Label 1	Label 0
Adult Census [15] Income	48,842	14	Sex-Male Race-White	Sex-Female Race-Non White	High Income	Low Income
Compas Software [39]	7,214	28	Sex-Male Race-Caucasian	Sex-Female Race-Non Caucasian	Did not Reoffend	Reoffend
Bank Marketing [14]	45,211	17	Age-Young	Age-Old	Subscriber	Non-subscriber
Default Credit [16]	13,636	23	Sex-Male	Sex-Female	Default	Not Default
Medical Expenditure (MEPS16) [18]	15,675	138	Race-White	Race-Non White	Utilized Benefits	Not Utilized Benefits

Table 3: Fairness and model utility metrics of the unrepaired (original) DNN.

Dataset	Unfair Neural Network					
	Validation			Test		
	EOD	F1	Accuracy	EOD	F1	Accuracy
Adult (Sex)	9.704% \pm 1.405	0.68 \pm 0.005	0.857 \pm 0.003	11.639% \pm 2.326	0.667 \pm 0.008	0.851 \pm 0.003
Adult (Race)	8.352% \pm 2.806			8.251% \pm 3.195		
COMPAS (Sex)	2.307% \pm 0.679	0.968 \pm 0.002	0.97 \pm 0.002	2.522% \pm 0.817		
COMPAS (Race)	2.022% \pm 0.884			2.96% \pm 1.088		
Bank	13.752% \pm 2.567	0.553 \pm 0.006	0.842 \pm 0.003	14.665% \pm 2.114	0.553 \pm 0.004	0.84 \pm 0.003
Default	9.068% \pm 1.782	0.538 \pm 0.006	0.774 \pm 0.007	8.962% \pm 1.772	0.53 \pm 0.006	0.769 \pm 0.007
MEPS16	20.167% \pm 2.32	0.543 \pm 0.009	0.79 \pm 0.005	20.641% \pm 2.527	0.533 \pm 0.01	0.788 \pm 0.009

Table 4: Fairness and model utility metrics of the repaired DNN.

Dataset	Repaired Neural Network					
	Simulated Annealing			Random Walk		
	EOD	F1	Accuracy	EOD	F1	Accuracy
Adult (Sex)	5.618% \pm 0.807	0.668 \pm 0.005	0.856 \pm 0.004	7.250% \pm 1.697	0.652 \pm 0.01	0.849 \pm 0.004
Adult (Race)	3.298% \pm 2.069	0.667 \pm 0.005	0.854 \pm 0.003	4.976% \pm 1.816	0.656 \pm 0.008	0.849 \pm 0.004
COMPAS (Sex)	0.468% \pm 0.548	0.955 \pm 0.004	0.959 \pm 0.003	2.921% \pm 1.446	0.954 \pm 0.08	0.957 \pm 0.008
COMPAS (Race)	0.56% \pm 0.71	0.955 \pm 0.002	0.959 \pm 0.002	2.239% \pm 1.003	0.954 \pm 0.005	0.957 \pm 0.004
Bank	0.871% \pm 0.509	0.547 \pm 0.007	0.892 \pm 0.01	7.257% \pm 3.533	0.537 \pm 0.014	0.888 \pm 0.01
Default	1.14% \pm 0.98	0.529 \pm 0.007	0.794 \pm 0.016	2.749% \pm 0.827	0.519 \pm 0.006	0.79 \pm 0.015
MEPS16	4.589% \pm 1.294	0.535 \pm 0.01	0.86 \pm 0.006	8.426% \pm 2.311	0.507 \pm 0.02	0.853 \pm 0.005

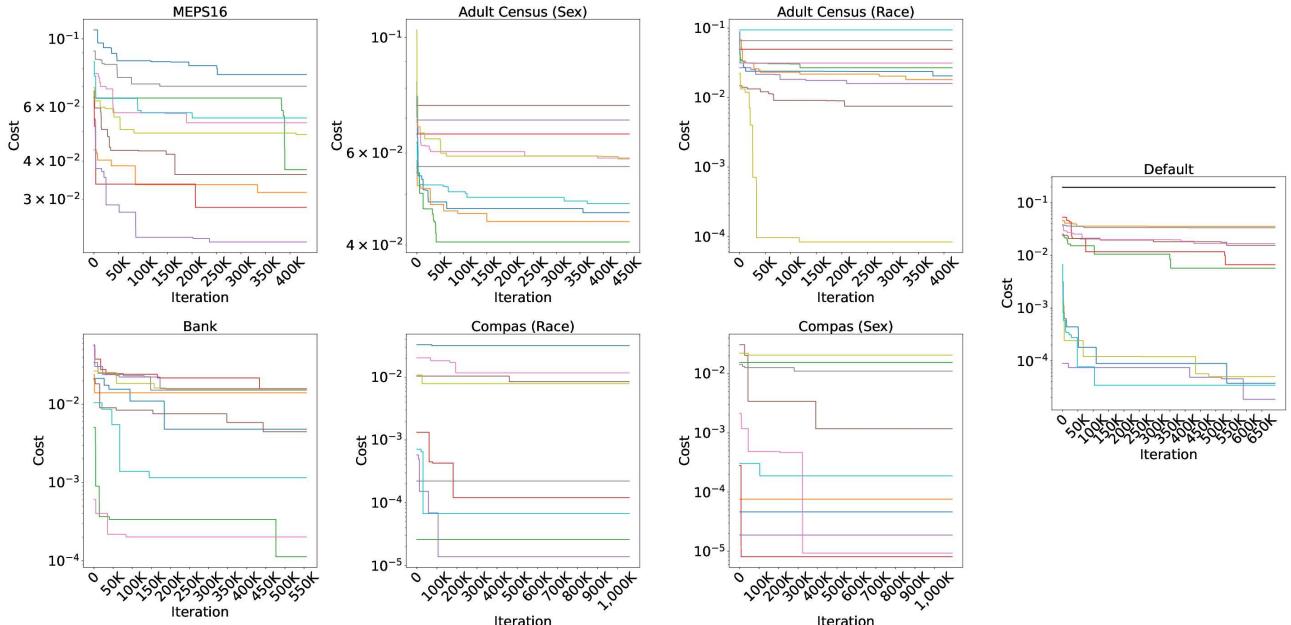


Figure 3: Evolution of the cost of the best state (logarithmic scale) found during the search on validation dataset using Simulated Annealing (SA) with 7 benchmarks and 10 seeds each.

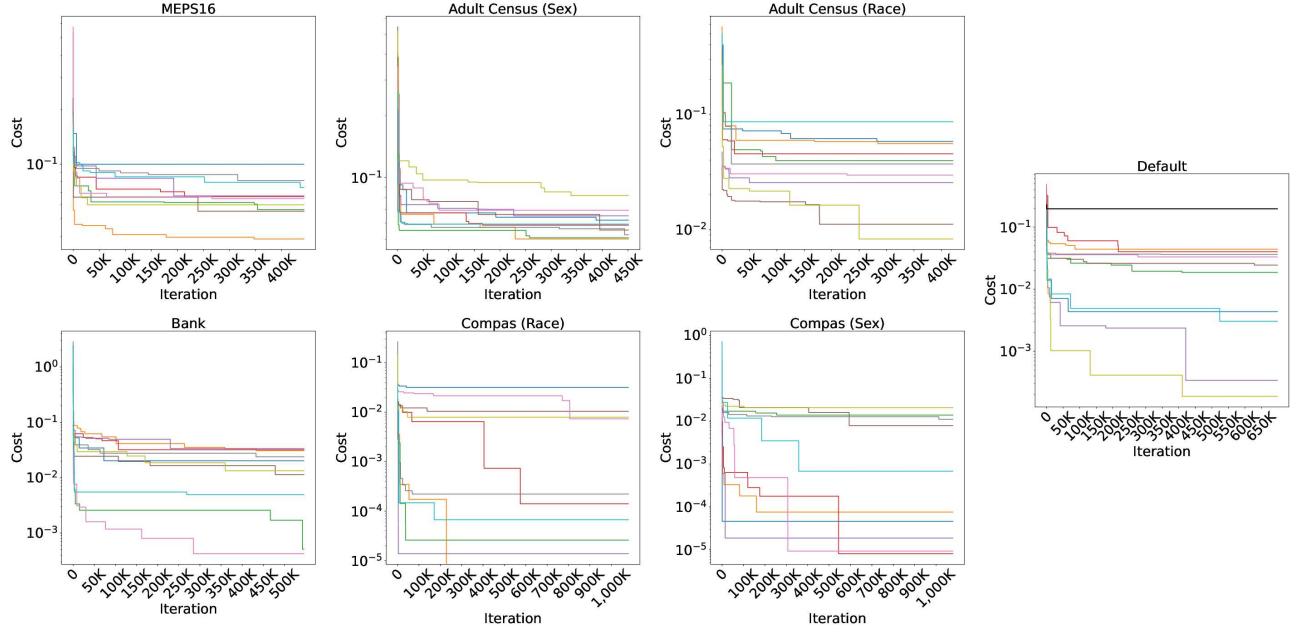


Figure 4: Evolution of the cost of the best state (logarithmic scale) found during the search on validation dataset using Random Walk (RW) with 7 benchmarks and 10 seeds each.

effective algorithm in identifying the desirable state during the search procedure, subject to the F1 threshold. On the test set, we observe that both NEUFAIR algorithms improve the EOD. SA outperforms RW for 4 datasets (*Adult (Sex)*, *Bank*, *Default*, and *MEPS16*). We observe the largest absolute improvement of 12.215% in EOD (20.641% to 8.426%) for the *MEPS16* dataset using SA. The largest relative improvement of 69.43% (8.962% to 2.749%) was observed in the *Default* dataset using SA. However, SA did not improve the EOD on the *COMPAS (Sex)*, which is a pathological case.

Table 5: Fairness metrics for COMPAS (Sex)

Model	EOD for 10 seeds									
	2.2%	2.5%	2.35%	1.87%	4.1%	2.15%	1.38%	4.26%	0.79%	3.6%
Unfair	2.2%	2.5%	2.35%	1.87%	4.1%	2.15%	1.38%	4.26%	0.79%	3.6%
Repaired (SA)	1.97%	4.28%	2.35%	4.58%	4.37%	0.12%	0.47%	4.83%	5.5%	0.75%
Repaired (RW)	1.25%	3.61%	3.19%	2.5%	2.4%	0.53%	0.28%	4.54%	3.07%	0.92%

Pathological Case. Table 5 highlights the individual EOD scores for all seeds of *COMPAS (Sex)* with the original and repaired networks. Here, we observe that SA improves the EOD in 4 cases, does not improve in 1 case, and has a worse EOD in 5 cases. RW improves the EOD in 5 cases and has a worse EOD in the other 5 cases. There are two reasons for the different trends in *COMPAS (Sex)*. First, there is a discrepancy in the original model statistics between the validation and test splits. The difference arises because the DNN tends to overfit slightly on the training and validation sets. This is evident from model metrics in Table 3. Second, the *COMPAS* datasets by default are quite fair (according to EOD) compared to other datasets as their unfairness is around 2%. During the search, the SA algorithm correctly identifies a desirable state according to the validation dataset and performs better than RW. However, given the low initial unfairness and the model's tendency to perform better

on the training and validation sets, we observe minor discrepancies in the test set. The average unfairness, however, increases by 0.399% (2.522% to 2.921%), which is negligible. For both NEUFAIR algorithms, the validation fairness improves in the *COMPAS (Sex)* experiment.

Answer RQ1: The randomized algorithms effectively mitigate unfairness via dropouts by de-activating a desirable subset of neurons. On the test set, Fairness improves by up to 69%. The SA algorithm performs better than the RW algorithm.

5.3.2 Dropout strategies improving both fairness and utility (RQ2). The second research question is answered by observing the F1 score and accuracy (model utility metrics), besides the EOD score (fairness) in Table 3 vs. Table 4. For all datasets, the F1 score decreases for the validation and test sets. The decrease in F1 score is accompanied by an improvement in fairness which highlights the tradeoff between the model utility and fairness. However, we observe that the accuracy increases for the *Bank*, *Default*, and *MEPS16* datasets. The biggest improvement in accuracy is observed in *MEPS16*, where the validation and test accuracy increase from 0.79 and 0.788 to 0.86 and 0.853 for SA and to 0.856 to 0.851 for RW. For *Default*, the validation and test accuracies increase from 0.774 and 0.769 to 0.794 and 0.79 for both SA and RW. For *Bank*, the validation and test accuracies increase from 0.842 and 0.84 to 0.892 and 0.88 for SA and 0.882 and 0.881 for RW.

The opposing trends of F1 score and accuracy can be attributed to the increase in negative predictions (0 is the negative class and 1 is positive) as neurons continue to be dropped out from the DNN. The increase in negative predictions favors the true negatives as the datasets are highly imbalanced with more negative samples. The *Default*, *Bank*, and *MEPS16* datasets have 78%, 88%, and 83% of

the data belonging to the negative class. The trained DNN models were optimized for F1 score on the validation dataset to improve performance on the underrepresented positive class. The accuracy increases as the increase in true negatives far outweighs the drop in true positives. The F1 score, however, always decreases as we lose precision and recall when the true positives decrease and the false negatives increase.

Answer RQ2: The overall accuracy, as a model utility metric, may improve along with fairness using dropout strategies, as accuracy does not account for false positives and false negatives in imbalanced datasets. However, the F1 score model utility metric decreases as fairness increases.

5.3.3 Hyperparameters of randomized algorithms and fairness (RQ3). The NEUFAIR algorithms have 4 hyperparameters: F1 threshold multiplier, the minimum and maximum number of neurons to drop, time-out limit, and F1 penalty multiplier.

- The F1 threshold multiplier is inversely proportional to fairness improvement. For example, as highlighted in RQ1, the fairness of the *Default* experiment improves when the F1 threshold is reduced from 0.98 to 0.96.
- Decreasing the minimum number of neurons and increasing the maximum number of neurons can have a positive effect on fairness, provided the time-out limit increases. Consider two ranges $[n_{l1}, n_{u1}]$ and $[n_{l2}, n_{u2}]$ such that $n_{l1} < n_{l2} < n_{u2} < n_{u1}$. The search space of $[n_{l2}, n_{u2}]$ is then a sub-space of $[n_{l1}, n_{u1}]$.
- The time-out positively affects fairness as the NEUFAIR algorithms have more time to explore the search space.
- The F1 penalty multiplier has a more nuanced effect on the unfairness as it controls the exploration vs. exploitation trade-offs. A low penalty multiplier increases the probability of state transitions when the F1 score is less than our threshold, increasing the SA run's randomness to explore more of the search space outside of the current best state. A high penalty multiplier would decrease the probability of state transitions, thereby encouraging the SA algorithm to exploit the search space near the current best state. The penalty multiplier p directly affects the cost difference ΔE between states with acceptable and unacceptable F1 scores. From Line 14 in Algorithm 2, we can see that a higher cost difference results in a lower transition probability. For a fixed temperature, a higher p would increase the cost difference between good and bad states, thereby decreasing the probability of transitioning and encouraging exploitation. A lower p would result in a lower cost difference and increase the transition probability, thereby encouraging exploration. To determine the effect of the penalty multiplier, we run experiments with the *Adult*, *MEPS16*, *Default*, and *Bank* datasets with one seed by varying the F1 penalty multiplier $p \in \{0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$ and F1 threshold multiplier to 0.98. Figure 5 shows the results. The red dot in the *Default* dataset corresponds to a run in which SA failed to find a desirable state within the F1 threshold. We can see that *Adult* dataset performs better on lower F1 penalties. The *Bank* dataset has a “sweet” spot in between while *Default* performs worse on intermediate F1 penalties. The *MEPS16* does not follow a clear trend. The results show that the F1 penalty multiplier is a

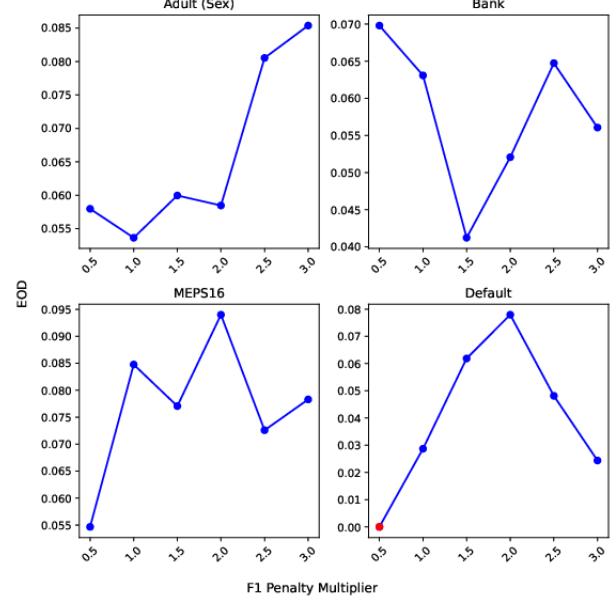


Figure 5: Effect of F1 penalty multiplier on the EOD

hyperparameter that trades off explorations vs. exploitation and requires tuning for each dataset and model.

Answer RQ3: The fairness of the repaired model can improve as the F1 threshold decreases, the number of neurons to drop increases, and the time-out limit increases for both NEUFAIR-SA and NEUFAIR-RW algorithms. The F1 penalty multiplier controls a trade-off between explorations vs. exploitation, which requires tuning for each specific benchmark. A higher F1 penalty multiplier encourages SA to exploit the local solutions, whereas a lower penalty encourages exploring the global space.

5.3.4 Comparing to the state-of-the-art (RQ4). We compare the efficacy of NEUFAIR against a state-of-the-art post-processing bias mitigator. DICE [36] uses a fault localization technique via do logic to pinpoint a single neuron that significantly influences fairness. The approach computes the amounts of variation in the outcomes of each layer that depends on the sensitive attributes and de-activates a single neuron with the highest effects on fairness. Since the search space evaluates a single neuron, we brute-force the entire search space of one neuron dropout in linear time. We note that the best improvement, feasible by DICE, cannot be more than the brute-force search. Table 6 highlights the test fairness after dropping a single neuron that most impacts fairness. The results show that NEUFAIR substantially outperforms DICE in improving fairness. We also note a similar trend with COMPAS (Sex), where the unfairness increases slightly. The largest discrepancy in fairness improvement between NEUFAIR and DICE is observed in the *MEPS16* dataset, with an absolute difference of 10.778% and a relative improvement of 56.124% (19.204% and 8.426%) using SA.

Table 6: Fairness and Utility of DICE [36].

Dataset	Test Dataset		
	EOD	F1	Accuracy
Adult (Sex)	$10.453\% \pm 2.266$	0.658 ± 0.007	0.851 ± 0.003
Adult (Race)	$8.092\% \pm 3.253$	0.662 ± 0.008	0.851 ± 0.003
COMPAS (Sex)	$3.229\% \pm 0.774$	0.964 ± 0.005	0.957 ± 0.004
COMPAS (Race)	$2.964\% \pm 1.088$	0.965 ± 0.003	0.968 ± 0.003
Bank	$12.205\% \pm 2.731$	0.567 ± 0.005	0.862 ± 0.01
Default	$5.845\% \pm 1.816$	0.528 ± 0.004	0.791 ± 0.007
MEPS16	$19.204\% \pm 2.592$	0.54 ± 0.012	0.8 ± 0.009

Table 7: Simulated Annealing compared to Brute-force search strategies.

Seed	Unfair DNN EOD	Repaired Neural Network EOD		
		SA	Brute Force	EOD Delta
1	14.087%	6.441%	6.415%	0.026
2	8.874%	5.306%	5.306%	0.0
3	8.352%	5.755%	5.755%	0.0
4	17.896%	13.823%	13.743%	0.08
5	8.062%	5.304%	5.25%	0.054

Table 8: Distribution of Best, Good, and Bad states in the search space.

Seed	Best State		Good State		Bad State	
	Count	Likelihood	Count	Likelihood	Count	Likelihood
1	1	2.3×10^{-8}	43,789	0.001	41,669,509	0.967
2	11	2.5×10^{-7}	6,717	1.0×10^{-4}	40,062,436	0.93
3	8	1.8×10^{-7}	1,795	4.1×10^{-5}	42,997,093	0.998
4	1	2.3×10^{-8}	790	1.8×10^{-5}	42,467,272	0.985
5	1	2.3×10^{-8}	13,623	3.16×10^{-4}	40,197,630	0.933

Answer RQ4: NEUFAIR outperforms the state-of-the-art unfairness mitigation algorithm DICE [36] across all benchmarks. The relative improvement of NEUFAIR is 56% higher than DICE.

6 Discussion

Effectiveness of Randomized Algorithms. To understand why randomized algorithms are effective in mitigating unfairness, we compare them to brute-force strategies. We create a small DNN for the *Adult (Sex)* dataset with 2 hidden layers, each with 16 neurons, to allow the brute force search to explore the entire space. We use a restricted search space with $n_l = 4$ and $n_u = 9$ that has a total of 43,076,484 states. The time-out for SA runs is one hour. The brute force exhaustive search takes 60 hours. Table 7 highlights the results of both search strategies. SA finds the global optimal solution for two seeds in the search space. The EOD Delta column highlights the difference in cost of the states identified by SA and brute force. The highest cost difference of 0.054 shows the efficacy of randomized algorithms in identifying good states. SA is faster than brute force in reaching a good state as it follows a cost gradient to choose a path to a good state by always accepting states that improve the cost. Additionally, SA uses a probabilistic transition function that can transition to worse states and potentially escape local minima.

Gradient-based methods often get stuck in local minima. Table 8 shows the distribution of the *Best*, *Good*, and *Bad* states. The best states are the states with the global optimal cost. The cost of the good states is within 0.05 of the optimal cost. Bad states are states whose F1 scores are less than the 98% of the baseline F1 score. We observe that the best and good states occupy a negligible amount of the search space. On the other hand, over 90% of the states are bad states. This suggests that the good states occupy small “pockets” of the search space. The cost gradient effectively leads SA to identify such pockets and the probabilistic transitions prevent it from getting stuck in local minima.

Limitations. Since we used randomized algorithms, our results might be sub-optimal. While SA provides statistical guarantees on the confidence and running time to reach an optimal state, in practice such guarantees might require multiple runs of algorithms over a long period of time. Additionally, we observe a discrepancy between the validation and test fairness improvements where the test set improves less than the validation set in both NEUFAIR algorithms. For our experiments, we use Equalized Odds Difference (EOD) only. However, the cost function can be modified accordingly with any fairness metric.

Threat to Validity. To address the internal validity and ensure our finding does not lead to an invalid conclusion, we follow the established SE guidelines and repeat the experiments 10 times, reporting both the average and 95% confidence intervals. In addition, we study the progress of algorithms during their runs, not just the outcomes. To ensure that our results are generalizable and address external validity, we perform our experiments on seven DNN benchmarks with various architectures. While similar architectures have been used in the prior works [36, 54, 57, 58], it is an open problem whether these datasets and DNN models are sufficiently representative for showing the effectiveness of NEUFAIR.

7 Related Work

Since the main focus of this paper is unfairness mitigation, we primarily focus on the prior work on mitigation.

A) Pre-processing Techniques exploit the space of input data to mitigate fairness defects [17, 20, 29, 51]. Reweighting [29] augments the dataset where the data points in each group-label combination are weighted differently to ensure fairness. The reweighting process minimizes the discrepancy between the observed and expected probabilities of the favorable outcome occurring with the sensitive attribute. FAIRMASK [37] uses the fact that protected attributes can be inferred by some combinations of other non-protected attributes (the “proxy” problem) to reduce the influence of protected attributes in the inference stage. FAIR-SMOTE [8] used under-sampling and over-sampling techniques [11] to balance data based on class and sensitive attributes. MAAT [12] proposes an ensemble training approach to mitigating unfairness. They propose a framework to train fairness and performance ML models and combine their outputs to make predictions. The performance model is identical to traditional ML algorithms and remains unchanged. For the fairness model, they design a debugging strategy based on prior work [9, 48] that attribute bias in training data to selection bias and label bias. LTDD [33] hypothesizes that biased features in training data contribute to unfairness. They develop a linear regression-based algorithm that

measures the association between non-sensitive and sensitive features in the dataset to obtain unbiased features. The non-sensitive features are then modified to remove the biased parts to become independent of the sensitive features. The sensitive features are dropped from the training dataset after debiasing.

The works above tamper with the precious training data samples that might lead to unrealistic representations of different protected groups. Additionally, it may not always be feasible to modify training data. Our work is tangential, focusing on post-processing mitigation that does not modify or access the training dataset and works on pre-trained unfair DNNs.

B) In-processing techniques improve fairness during training. *Adversarial debiasing* [52], employs adversarial learning to develop a classifier that hinders the ability of adversaries to determine sensitive attributes from ML predictions. Another approach, the *prejudice remover* [30], incorporates a fairness-centric adjustment into the loss function to balance accuracy and fairness. These strategies necessitate alterations in either the loss function or the model parameters. Seldonian [26, 34, 45] presented a technique that allows users to directly specify (arbitrary) undesirable behaviors as constraints and enforce them during training. FAIRWAY [10] combines pre-processing and in-processing mitigation techniques to improve fairness. PARFAIT-ML [46] is a gray-box evolutionary search algorithm that explores the ML hyperparameters to find configurations that minimize fairness while maintaining an acceptable accuracy. While mitigating unfairness during training might be less intrusive than changing the training datasets, it often requires changing the training algorithms. This might only be effective for a particular algorithm and fairness definition. Additionally, the training process is often expensive and might not be feasible for existing systems.

C) Post-processing techniques aim to modify the prediction outcomes of ML models to reduce discrimination [25, 30, 38, 44]. Equalized Odds Processing (EOP) [25] modifies the output of a biased binary classifier to improve fairness. EOP solves an optimization problem using a linear program that utilizes the protected attribute, predictions, and true labels from a dataset to create an unbiased predictor. NEUFAIR, however, does not require training auxiliary models to mitigate unfairness. Additionally, NEUFAIR does not rely on randomization for fairness. Predictions during inference time are deterministic after the neurons are dropped from the DNN. Woodworth et al. [49] also show that EOP requires the biased classifier to be Bayes optimal, which is practically impossible to learn from finite samples of data. Zhang and Sun [55] adaptively intervened in the input data features and DNN internal neurons to improve fairness. CARE [42] presented a Particle Swarm Optimization (PSO) algorithm to repair DNNs. FAIRE [32] proposes a method to repair unfair DNNs by altering the activations of neurons. Inspired by program repair techniques, FAIRE first groups neurons representing protected and non-protected features in the dataset into two categories. To create the two categories, FAIRE creates a clone of the unfair model and retrains it to predict the protected attribute. By comparing the activations of the corresponding neurons in both networks, FAIRE determines which neurons represent the protected and non-protected features. Finally, the protected neuron activations are penalized, and the non-protected neuron activations are promoted. NEUFAIR, however, does not retrain the model and does

not separately use the protected attribute to identify unfair neurons. We find an unfair subset of neurons given the entire feature vector, thereby accounting for the causal influence of the protected features on the non-protected features in the training dataset. Our analysis considers a group of neurons together rather than identifying the contributions of individual neurons. FAIRNEURON [21] proposes an algorithm to mitigate unfairness by selectively retraining neurons in a DNN. On a high level, the algorithm identifies input-output neuron paths corresponding to each sample in the training dataset that induces bias in the prediction. Then, the dataset is split into biased and unbiased subsets corresponding to the neuron activations in the paths. Finally, the algorithm retrains the model with both training data splits by enabling random dropouts in the biased paths and disabling any dropouts in the unbiased paths. While FAIRNEURON uses dropout to improve fairness, it fundamentally differs from NEUFAIR. We do not retrain the DNN by using the training dataset. Also, NEUFAIR utilizes deterministic dropout during inference time to identify a desirable subset of neurons to drop that has a minimal impact on F1 score but the maximal impact on fairness, rather than dropping neurons randomly.

8 Conclusion

In this paper, we tackle the problem of mitigating unfairness in pre-trained DNNs using the dropout method. We showed that the neural dropout problem over the DNN models is computationally hard and presented NEUFAIR, a family of randomized algorithms to efficiently and effectively improve the fairness of DNNs. Our experiments showed that NEUFAIR can identify an ideal subset of neurons to drop that disproportionately contribute to unfairness (leading to up to 69% fairness improvement) and outperform a state-of-the-art post-processing bias mitigator.

For future work, there are a few exciting directions. First, we can analyze the top k states found during the run instead of choosing the best state. Such a strategy could reduce the difference in improvement observed between the validation and test sets. We can extend to multi-valued protected attributes (e.g., age groups) or optimize for more than one protected attribute at a time (e.g., race and sex). Second, we plan to leverage dropouts in the fine-tuning process to understand the effect of randomized algorithms on mitigating the unfairness and toxicity of large language models. Lastly, alternate strategies like genetic algorithms may be used to solve the combinatorial optimization problem of neuron dropout. We use randomized algorithms for their simplicity in efficiently solving combinatorial optimization problems. The primary difference between SA/RW and genetic algorithms is that genetic algorithms create a population of candidate solutions at each step instead of a single solution. This approach can be promising as it enables us to explore more states of the search space simultaneously.

Data Availability

Our open-source tool NEUFAIR with all experimental subjects are publicly accessible on Zenodo [13] and [GitHub](https://github.com).

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. CNS-2230060 and CNS-2230061.

References

[1] Alekh Agarwal, Alina Beygelzimer, Miroslav Dudik, John Langford, and Hanna Wallach. 2018. A reductions approach to fair classification. In *International Conference on Machine Learning*. PMLR, 60–69.

[2] Aniya Aggarwal, Pranay Lohia, Seema Nagar, Kuntal Dey, and Diptikalyan Saha. 2019. Black Box Fairness Testing of Machine Learning Models. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2019)*. 625–635. <https://doi.org/10.1145/3338906.3338937>

[3] Walid Ben-Ameur. 2004. Computing the Initial Temperature of Simulated Annealing. *Computational Optimization and Applications* 29, 3 (01 Dec 2004), 369–385. <https://doi.org/10.1023/B:COAP.0000044187.23143.bd>

[4] Dimitris Bertsimas and John Tsitsiklis. 1993. Simulated annealing. *Statist. Sci.* 8, 1 (1993), 10–15.

[5] Sumon Biswas and Hridesh Rajan. 2021. Fair Preprocessing: Towards Understanding Compositional Fairness of Data Transformers in Machine Learning Pipeline. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Athens, Greece) (ESEC/FSE 2021)*. Association for Computing Machinery, New York, NY, USA, 981–993. <https://doi.org/10.1145/3468264.3468536>

[6] Ruth G Blumrosen. 1978. Wage discrimination, job segregation, and the title vii of the civil rights act of 1964. *U. Mich. Jl. Reform* 12 (1978), 397.

[7] Yuriy Brun and Alexandra Meliou. 2018. Software Fairness (ESEC/FSE 2018). 754–759. <https://doi.org/10.1145/3236024.3264838>

[8] Joymallya Chakraborty, Sudodeep Majumder, and Tim Menzies. 2021. Bias in Machine Learning Software: Why? How? What to Do?. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Athens, Greece) (ESEC/FSE 2021)*. Association for Computing Machinery, New York, NY, USA, 429–440. <https://doi.org/10.1145/3468264.3468537>

[9] Joymallya Chakraborty, Sudodeep Majumder, and Tim Menzies. 2021. Bias in machine learning software: why? how? what to do?. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Athens, Greece) (ESEC/FSE 2021)*. Association for Computing Machinery, New York, NY, USA, 429–440. <https://doi.org/10.1145/3468264.3468537>

[10] Joymallya Chakraborty, Sudodeep Majumder, Zhe Yu, and Tim Menzies. 2020. Fairway: a way to build fair ML software. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 654–665.

[11] Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. 2002. SMOTE: Synthetic Minority over-Sampling Technique. *J. Artif. Int. Res.* 16, 1 (jun 2002), 321–357.

[12] Zhenpeng Chen, Jie M. Zhang, Federica Sarro, and Mark Harman. 2022. MAAT: a novel ensemble approach to addressing fairness and performance bugs for machine learning software. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (<conf-loc>, <city>Singapore</city>, <country>Singapore</country>, </conf-loc>) (ESEC/FSE 2022)*. Association for Computing Machinery, New York, NY, USA, 1122–1134. <https://doi.org/10.1145/3540250.3549093>

[13] Vishnu Asutosh Dasu. 2024. NeuFair: Neural Network Fairness Repair with Dropout. <https://doi.org/10.5281/zenodo.1266209>

[14] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <https://archive.ics.uci.edu/ml/datasets/bank+marketing>

[15] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. <https://archive.ics.uci.edu/ml/datasets/census+income>

[16] Dheeru Dua and Casey Graff. 2017. UCI Machine Learning Repository. [https://archive.ics.uci.edu/ml/datasets/statlog+\(german+credit+data\)](https://archive.ics.uci.edu/ml/datasets/statlog+(german+credit+data))

[17] Michael Feldman, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. 2015. Certifying and Removing Disparate Impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Sydney, NSW, Australia) (KDD '15). Association for Computing Machinery, New York, NY, USA, 259–268. <https://doi.org/10.1145/2783258.2783311>

[18] Agency for Healthcare Research and Quality. [Online]. Medical Expenditure Panel Survey. <https://meps.ahrq.gov/mepsweb/>.

[19] Sainyam Galhotra, Yuriy Brun, and Alexandra Meliou. 2017. Fairness Testing: Testing Software for Discrimination (ESEC/FSE 2017). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3106237.3106277>

[20] Sainyam Galhotra, Karthikeyan Shanmugam, Prasanna Sattigeri, and Kush R Varshney. 2022. Causal feature selection for algorithmic fairness. In *Proceedings of the 2022 International Conference on Management of Data*. 276–285.

[21] Xuanqi Gao, Juan Zhai, Shiqing Ma, Chao Shen, Yufei Chen, and Qian Wang. 2022. FairNeuron: Improving Deep Neural Network Fairness with Adversary Games on Selective Neurons. In *Proceedings of the 44th International Conference on Software Engineering* (Pittsburgh, Pennsylvania) (ICSE '22). Association for Computing Machinery, New York, NY, USA, 921–933. <https://doi.org/10.1145/3497863.3498863>

[22] Usman Gohar, Sumon Biswas, and Hridesh Rajan. 2023. Towards Understanding Fairness and its Composition in Ensemble Machine Learning. [arXiv:2212.04593 \[cs.LG\]](https://arxiv.org/abs/2212.04593)

[23] Home Credit Group. 2018. Home Credit Default Risk. <https://www.kaggle.com/competitions/home-credit-default-risk/overview>.

[24] Bruce Hajek. 1988. Cooling schedules for optimal annealing. *Mathematics of operations research* 13, 2 (1988), 311–329.

[25] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of Opportunity in Supervised Learning. In *NIPS*.

[26] Austin Hoag, James E. Kostas, Bruno Castro da Silva, Philip S. Thomas, and Yuriy Brun. 2023. Seldonian Toolkit: Building Software with Safe and Fair Machine Learning. *ICSE (Demo Track)* (2023).

[27] David Ingold and Spencer Soper. 2016. Amazon Doesn't Consider the Race of Its Customers. Should It? <https://www.bloomberg.com/graphics/2016-amazon-same-day/>. Online.

[28] Surya Mattu Julia Angwin, Jeff Larson and Lauren Kirchner. 2021. Machine Bias. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>. Online.

[29] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and information systems* 33, 1 (2012), 1–33.

[30] Faisal Kamiran, Asim Karim, and Xiangliang Zhang. 2012. Decision Theory for Discrimination-Aware Classification. In *2012 IEEE 12th International Conference on Data Mining*. 924–929. <https://doi.org/10.1109/ICDM.2012.45>

[31] Arjun Kharpal. 2018. TECH Health care start-up says A.I. can diagnose patients better than humans can, doctors call that 'dubious'. *CNBC* (2018). <https://www.cnbc.com/2018/06/28/babylon-claims-its-ai-can-diagnose-patients-better-than-doctors.html>

[32] Tianlin Li, Xiaofei Xie, Jian Wang, Qing Guo, Aishan Liu, Lei Ma, and Yang Liu. 2023. Faire: Repairing Fairness of Neural Networks via Neuron Condition Synthesis. *ACM Trans. Softw. Eng. Methodol.* 33, 1, Article 21 (nov 2023), 24 pages. <https://doi.org/10.1145/3617168>

[33] Yanhui Li, Linghan Meng, Lin Chen, Li Yu, Di Wu, Yuming Zhou, and Baowen Xu. 2022. Training data debugging for the fairness of machine learning software. In *Proceedings of the 44th International Conference on Software Engineering* (Pittsburgh, Pennsylvania) (ICSE '22). Association for Computing Machinery, New York, NY, USA, 2215–2227. <https://doi.org/10.1145/3510003.3510091>

[34] Blossom Metevier, Stephen Giguere, Sarah Brockman, Ari Kobren, Yuriy Brun, Emma Brunskill, and Philip S Thomas. 2019. Offline contextual bandits with high probability fairness guarantees. *Advances in neural information processing systems* 32 (2019).

[35] Debasis Mitra, Fabio Romeo, and Alberto Sangiovanni-Vincentelli. 1986. Convergence and finite-time behavior of simulated annealing. *Advances in applied probability* 18, 3 (1986), 747–771.

[36] Verya Monjezi, Ashutosh Trivedi, Gang Tan, and Saeid Tizpaz-Niari. 2023. Information-Theoretic Testing and Debugging of Fairness Defects in Deep Neural Networks. In *Proceedings of the 45th International Conference on Software Engineering* (Melbourne, Victoria, Australia) (ICSE '23). IEEE Press, 1571–1582. <https://doi.org/10.1109/ICSE48619.2023.00136>

[37] Kewen Peng, Joymallya Chakraborty, and Tim Menzies. 2022. FairMask: Better Fairness via Model-Based Rebalancing of Protected Attributes. *IEEE Trans. Softw. Eng.* 49, 4 (nov 2022), 2426–2439. <https://doi.org/10.1109/TSE.2022.3220713>

[38] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q. Weinberger. 2017. On Fairness and Calibration (*NIPS'17*). 5684–5693.

[39] ProPublica. 2021. Compas Software Ananlysis. <https://github.com/propublica/compas-analysis>. Online.

[40] Nitish Srivastava. 2013. Improving neural networks with dropout. *University of Toronto* 182, 566 (2013), 7.

[41] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *J. Mach. Learn. Res.* 15, 1 (jan 2014), 1929–1958.

[42] Bing Sun, Jun Sun, Long H Pham, and Jie Shi. 2022. Causality-based neural network repair. In *Proceedings of the 44th International Conference on Software Engineering*. 338–349.

[43] Jim Tankersley. 2023. Black Americans Are Much More Likely to Face Tax Audits, Study Finds. *The New York Times* (31 Jan 2023). <https://www.nytimes.com/2023/01/31/us/politics/black-americans-irs-tax-audits.html>

[44] Guanhong Tao, Weisong Sun, Tingxu Han, Chunrong Fang, and Xiangyu Zhang. 2022. RULER: discriminative and iterative adversarial training for deep neural network fairness. In *Proceedings of the 30th acm joint european software engineering conference and symposium on the foundations of software engineering*. 1173–1184.

[45] Philip S. Thomas, Bruno Castro da Silva, Andrew G. Barto, Stephen Giguere, Yuriy Brun, and Emma Brunskill. 2019. Preventing undesirable behavior of intelligent machines. *Science* 366, 6468 (2019), 999–1004. <https://doi.org/10.1126/science.aag3311> arXiv:<https://www.science.org/doi/pdf/10.1126/science.aag3311>

[46] Saeid Tizpaz-Niari, Ashish Kumar, Gang Tan, and Ashutosh Trivedi. 2022. Fairness-aware configuration of machine learning libraries. In *Proceedings of the 44th International Conference on Software Engineering*. 909–920.

[47] Sakshi Udeshi, Pryanshu Arora, and Sudipta Chattopadhyay. 2018. Automated directed fairness testing. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*. 98–108.

[48] Michael Wick, swetasudha panda, and Jean-Baptiste Tristan. 2019. Unlocking Fairness: a Trade-off Revisited. In *Advances in Neural Information Processing Systems*. H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.), Vol. 32. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2019/file/373e4c5d8edfa8b74fd4b6791d0cf6dc-Paper.pdf

[49] Blake Woodworth, Suriya Gunasekar, Mesrob I. Ohannessian, and Nathan Srebro. 2017. Learning Non-Discriminatory Predictors. arXiv:1702.06081 [cs.LG]

[50] Normen Yu, Gang Tan, and Saeid Tizpaz-Niari. 2023. FairLay-ML: Intuitive Remedies for Unfairness in Data-Driven Social-Critical Algorithms. arXiv:2307.05029 [cs.LG]

[51] Richard Zemel, Yu Wu, Kevin Swersky, Toniam Pitassi, and Cynthia Dwork. 2013. Learning Fair Representations. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28* (Atlanta, GA, USA) (ICML '13). JMLR.org, III–325–III–333.

[52] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. 335–340.

[53] Jie M Zhang and Mark Harman. 2021. "Ignorance and Prejudice" in Software Fairness. In *2021 IEEE/ACM 43rd International Conference on Software Engineering* (ICSE). IEEE, 1436–1447.

[54] Lingfeng Zhang, Yueling Zhang, and Min Zhang. 2021. Efficient White-Box Fairness Testing through Gradient Search. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis (Virtual, Denmark) (ISSTA 2021)*. 103–114. <https://doi.org/10.1145/3460319.3464820>

[55] Mengdi Zhang and Jun Sun. 2022. Adaptive Fairness Improvement Based on Causality Analysis. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (Singapore, Singapore) (ESEC/FSE 2022)*. Association for Computing Machinery, New York, NY, USA, 6–17. <https://doi.org/10.1145/3540250.3549103>

[56] Peixin Zhang, Jingyi Wang, Jun Sun, Guoliang Dong, Xinyu Wang, Xingen Wang, Jin Song Dong, and Ting Dai. 2020. White-Box Fairness Testing through Adversarial Sampling (ICSE '20). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3377811.3380331>

[57] Peixin Zhang, Jingyi Wang, Jun Sun, Guoliang Dong, Xinyu Wang, Xingen Wang, Jin Song Dong, and Ting Dai. 2020. White-box fairness testing through adversarial sampling. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 949–960.

[58] Haibin Zheng, Zhiqing Chen, Tianyu Du, Xuhong Zhang, Yao Cheng, Shouling Ti, Jingyi Wang, Yue Yu, and Jinyin Chen. 2022. NeuronFair: Interpretable White-Box Fairness Testing through Biased Neuron Identification. In *2022 IEEE/ACM 44th International Conference on Software Engineering (ICSE)*. 1519–1531. <https://doi.org/10.1145/3510003.3510123>

Received 2024-04-12; accepted 2024-07-03