



Probabilistic Verification of Neural Networks Against Group Fairness

Bing Sun¹(✉), Jun Sun¹, Ting Dai², and Lijun Zhang^{3,4}

¹ Singapore Management University, Bras Basah, Singapore
bing.sun.2020@phdcs.smu.edu.sg

² Huawei, Changi, Singapore

³ SKLCS, University of Chinese Academy of Science, Beijing, China

⁴ Institute of Intelligent Software, Guangzhou, China

Abstract. Fairness is crucial for neural networks which are used in applications with important societal implication. Recently, there have been multiple attempts on improving fairness of neural networks, with a focus on fairness testing (e.g., generating individual discriminatory instances) and fairness training (e.g., enhancing fairness through augmented training). In this work, we propose an approach to formally verify neural networks against fairness, with a focus on independence-based fairness such as group fairness. Our method is built upon an approach for learning Markov Chains from a user-provided neural network (i.e., a feed-forward neural network or a recurrent neural network) which is guaranteed to facilitate sound analysis. The learned Markov Chain not only allows us to verify (with Probably Approximate Correctness guarantee) whether the neural network is fair or not, but also facilitates sensitivity analysis which helps to understand why fairness is violated. We demonstrate that with our analysis results, the neural weights can be optimized to improve fairness. Our approach has been evaluated with multiple models trained on benchmark datasets and the experiment results show that our approach is effective and efficient.

1 Introduction

In recent years, neural network based machine learning has found its way into various aspects of people's daily life, such as fraud detection [25], facial recognition [47], self-driving [13], and medical diagnosis [56]. Although neural networks have demonstrated astonishing performance in many applications, there are still concerns on their dependability. One desirable property of neural networks for applications with societal impact is fairness [2]. Since there are often societal biases in the training data, the resultant neural networks might be discriminative as well. This has been demonstrated in [53]. Fairness issues in neural networks are often more 'hidden' than those of traditional decision-making software programs since it is still an open problem on how to interpret neural networks.

Recently, researchers have established multiple formalization of fairness regarding different sub-populations [9, 21, 24, 28]. These sub-populations are often

determined by different values of protected features (e.g., race, religion and ethnic group), which are application-dependent. To name a few, group fairness requires that minority members should be classified at an approximately same rate as the majority members [9, 24], whereas individual discrimination (a.k.a. causal fairness) states that a machine learning model must output approximately the same predictions for instances which are the same except for certain protected features [21, 28]. We refer readers to [51] for detailed definitions of fairness. In this work, we focus on an important class of fairness called independence-based fairness, which includes the above-mentioned group fairness.

Recently, there have been multiple attempts on analyzing and improving fairness of neural networks, with a focus on fairness testing (e.g., generating individual discriminatory instances) and fairness training (e.g., enhancing fairness through augmented training). Multiple attempts [4, 27, 54, 58] have been made on testing machine learning models against individual discrimination, which aims to systematically generate instances that demonstrate individual discrimination. While these approaches have impressive performance in terms of generating such instances, they are incapable of verifying fairness. Another line of approaches is on fairness training [3, 12, 14, 16, 28, 36], this includes approaches which incorporate fairness as an objective in the model training phase [3, 12, 16], and approaches which adopt heuristics for learning fair classifiers [36]. While the experiment results show that these approaches improve fairness to certain extent, they do not guarantee that the resultant neural networks are fair.

In this work, we investigate the problem of verifying neural networks against independence-based fairness. Our aim is to design an approach which allows us to (1) show evidence that a neural network satisfies fairness if it is the case; (2) otherwise, provide insights on why fairness is not satisfied and how fairness can be potentially achieved; (3) provide a way of improving the fairness of the neural network. At a high-level, our approach is designed as follows. Given a neural network (i.e., either a feed-forward or recurrent neural network), we systematically sample behaviors of the neural network (e.g., input/output pairs), based on which we learn a Markov Chain model that approximates the neural network. Our algorithm guarantees that probabilistic analysis based on the learned Markov Chain model (such as probabilistic reachability analysis) is probably approximately correct (hereafter PAC-correct) with respect to any computational tree logic (CTL [11]) formulae. With the guarantee, we are thus able to verify fairness property of the neural network. There are two outcomes. One is that the neural network is proved to be fair, in which case the Markov Chain is presented as an evidence. Otherwise, sensitivity analysis based on the Markov Chain is carried out automatically. Such analysis helps us to understand why fairness is violated and provide hints on how the neural network could be improved to achieve fairness. Lastly, our approach optimizes the parameters of the ‘responsible’ neurons in the neural network and improve its fairness.

We have implemented our approach as a part of the SOCRATES framework [45]. We apply our approach to multiple neural network models (including feed-forward and recurrent neural networks) trained on benchmark datasets

which are the subject of previous studies on fairness testing. The experiment results show that our approach successfully verifies or falsifies all the models. It also confirms that fairness is a real concern and one of the networks (on the German Credit dataset) fails the fairness property badly. Through sensitivity analysis, our approach locates neurons which have the most contribution to the violation of fairness. Further experiments show that by optimizing the neural parameters (i.e., weights) based on the sensitivity analysis result, we can improve the model’s fairness significantly whilst keeping a high model accuracy.

The remaining of the paper is organized as follows. In Sect. 2, we review relevant background and define our problem. In Sect. 3, we present each step of our approach in detail. In Sect. 4, we evaluate our approach through multiple experiments. We review related work in Sect. 5 and conclude in Sect. 6. For appendix, please refer to the extended version of this paper [50].

2 Preliminary

In this section, we review relevant background and define our problem.

Fairness. For classification problems, a neural network N learns to predict a target variable O based on a set of input features X . We write Y as the prediction of the classifier. We further write $F \subseteq X$ as a set of features encoding some protected characteristics such as gender, age and race. Fairness constrains how N makes predictions. In the literature, there are multiple formal definitions of fairness [9, 21, 24, 28]. In this work, we focus on independence-based fairness, which is defined as follows.

Definition 1 (Independence-based Fairness (strict)). *A neural network N satisfies independence-based fairness (strict) if the protected feature F is statistically independent to the prediction Y . We write L as the prediction set and we have $\forall l \in L, \forall f_i, f_j \in F$ such that $i \neq j$,*

$$P(Y = l \mid F = f_i) = P(Y = l \mid F = f_j) \quad (1)$$

The definition states that, N ’s prediction is independent of the protected feature F . This definition is rather strict and thus unlikely to hold in practice. The following relaxes the above definition by introducing a positive tolerance ξ .

Definition 2 (Independence-based Fairness). *Let N be a neural network and ξ be a positive real-value constant. N satisfies independence-based fairness, with respect to ξ , if and only if, $\forall l \in L \forall f_i, f_j \in F$ such that $i \neq j$,*

$$|P(Y = l \mid F = f_i) - P(Y = l \mid F = f_j)| \leq \xi \quad (2)$$

Intuitively, the above definition states that N is fair as long as the probability difference is within the threshold ξ . In the following, we focus on Definition 2 as it is both more general and more practical compared to Definition 1.

Example 1. Let us take the network trained on the Census Income dataset [18] as an example. The dataset consists of 32k training instances, each of which contains 13 features. The task is to predict whether an individual’s income exceeds \$50K per year. An example instance x with a prediction y will be $x : \langle 3 \ 5 \ 3 \ 0 \ 2 \ 8 \ 3 \ 0 \ 1 \ 2 \ 0 \ 40 \ 0 \rangle$, $y : \langle 0 \rangle$. Note that all features are categorical (i.e., processed using binning). Among all features, gender, age and race are considered protected features. The model N trained based on the dataset is in the form of a six-layer fully-connected feed-forward neural network. The following is a fairness property defined based on the protected feature gender.

$$|P(Y = 1 \mid F = \textit{male}) - P(Y = 1 \mid F = \textit{female})| \leq 0.1 \quad (3)$$

Intuitively, the difference in the probability of predicting 1, for males and females, should be no more than 10%.

Our Problem. We are now ready to define our problem.

Definition 3 (The verification problem). *Let N be a neural network. Let ϕ be an independence-based fairness property (with respect to protected feature F and a threshold ξ). The fairness verification problem is to verify whether N satisfies ϕ or not.*

One way of solving the problem is through statistical model checking (such as hypothesis testing [40]). Such an approach is however not ideal. While it is possible to conclude whether N is fair or not (with certain level of statistical confidence), the result often provides no insight. In the latter case, we would often be interested in performing further analysis to answer questions such as whether certain feature or neuron at a hidden layer is particularly relevant to the fairness issue and how to improve the fairness. The above-mentioned approach offers little clue to such questions.

3 Our Approach

In this section, we present details of our approach. Our approach is shown in Algorithm 1. The first step is to learn a Markov Chain D which guarantees that probabilistic analysis such as probabilistic reachability analysis based on D is PAC-correct with respect to N . The second step is to apply probabilistic model checking [39] to verify D against the fairness property ϕ . In the third step, if the property ϕ is not verified, sensitivity analysis is performed on D which provides us information on how to improve N in terms of fairness. That is, we improve the fairness of the model by optimizing the neuron weights based on the sensitivity analysis results.

Note that our approach relies on building an approximation of the neural network in the form of Markov Chains. There are three reasons why constructing such an abstraction is beneficial. First, it allows us to reason about unbounded behaviors (in the case of a cyclic Markov Chains, which can be constructed from

recurrent neural networks as we show below) which are known to be beyond the capability of statistical model checking [40]. Second, the Markov Chain model allows us to perform analysis such as sensitivity analysis (e.g., to identify neurons responsible for violating fairness) as well as predict the effect of changing certain probability distribution (e.g., whether fairness will be improved), which are challenging for statistical methods. Lastly, in the case that the fairness is verified, the Markov Chain serves as a human-interpretable argument on why fairness is satisfied.

In the following, we introduce each step in detail. We fix a neural network N and a fairness property ϕ of the form $|P(Y = l | F = f_i) - P(Y = l | F = f_j)| \leq \xi$. We use the neural network trained on the Census Income dataset (refer to Example 1) as a running example.

Algorithm 1: <i>verify_repair</i> ($N, \phi, \mu\epsilon, \mu\delta$)	Algorithm 2: <i>learn</i> (N, S, ϵ, δ)
<ol style="list-style-type: none"> 1 Fix the set of states S; 2 Learn DTMC D by $\text{learn}(N, S, \frac{\mu\epsilon}{2}, 1 - \sqrt{1 - \mu\delta})$; 3 Estimate $P(Y = l F = f_i) \forall f_i \in F$; 4 Verify ϕ against ξ; if ϕ is verified then 5 return “Verified” and D; 6 else 7 Conduct sensitivity analysis on D; 8 Perform automatic repair of N; 9 return N'; 	<ol style="list-style-type: none"> 1 $W := 0$; 2 $A_W := 0$; 3 do 4 generate new sample trace ω 5 $W := W + \omega$; 6 update $A_W(p, q)$ for all $p \in S$ and $q \in S$; 7 update $H(n)$; 8 while $\exists p \in S, n_p < H(n)$ <p style="margin: 0;">Output: A_W</p>

3.1 Step 1: Learning a Markov Chain

In this step, we construct a Discrete-Time Markov Chain (DTMC) which approximates N (i.e., line 2 of Algorithm 1). DTMCs are widely used to model the behavior of stochastic systems [10], and they are often considered reasonably human-interpretable. Example DTMCs are shown in Fig. 1. The definition of DTMC is presented in [50]. Algorithm 2 shows the details of this step. The overall idea is to construct a DTMC, based on which we can perform various analysis such as verifying fairness. To make sure the analysis result on the DTMC applies to the original N , it is important that the DTMC is constructed in such a way that it preserves properties such as probabilistic reachability analysis (which is necessary for verifying fairness as we show later). Algorithm 2 is thus based on the recent work published in [10], which develops a sampling method for learning DTMC. To learn a DTMC which satisfies our requirements, we must answer three questions.

(1) *What are the states S in the DTMC?* The choice of S has certain consequences in our approach. First, it constrains the kind of properties that we are allowed to analyze based on the DTMC. As we aim to verify fairness, the states must minimally include states representing different protected features, and states representing prediction outcomes. The reason is that, with these

states, we can turn the problem of verifying fairness into probabilistic reachability analysis based on the DTMC, as we show in Sect. 3.2. What additionally are the states to be included depends on the level of details that we would like to have for subsequent analysis. For instance, we include states representing other features at the input layer, and states representing the status of hidden neurons. Having these additional states allows us to analyze the correlation between the states and the prediction outcome. For instance, having states representing a particular feature (or the status of a neuron of a hidden layer) allows us to check how sensitive the prediction outcome is with respect to the feature (or the status of a neuron). Second, the choice of states may have an impact on the cost of learning the DTMC. In general, the more states there are, the more expensive it is to learn the DTMC. In Sect. 4, we show empirically the impact of having different sizes of S . We remark that to represent continuous input features and hidden neural states using discrete states, we discretize their values (e.g., using binning or clustering methods such as *Kmeans* [42] based on a user-provided number of clusters). In our approach, for fairness verification only task, we include protected features and prediction outcomes. For sensitivity analysis task, we further include other features and hidden neural states.

(2) *How do we identify the transition probability matrix?* The answer is to repeatedly sample inputs (by sampling based on a prior probability distribution) and then monitor the trace of the inputs, i.e., the sequence of transitions triggered by the inputs. After sampling a sufficiently large number of inputs, the transition probability matrix then can be estimated based on the frequency of transitions between states in the traces. In general, the question of estimating the transition probability matrix of a DTMC is a well-studied topic and many approaches have been proposed, including frequency estimation, Laplace smoothing [10] and Good-Turing estimation [26]. In this work, we adopt the following simple and effective estimation method. Let W be a set of traces which can be regarded as a bag of transitions. We write n_p where $p \in S$ to denote the number transitions in W originated from state p . We write n_{pq} where $p \in S$ and $q \in S$ to be the number of transitions observed from state p to q in W . Let m be the total number of states in S . The transition probability matrix A_W (estimated based on W) is: $A_W(p, q) = \begin{cases} \frac{n_{pq}}{n_p} & \text{if } n_p \neq 0 \\ \frac{1}{m} & \text{otherwise} \end{cases}$. Intuitively, the probability of transition from state p to q is estimated as the number of transitions from p to q divided by the total number of transitions taken from state p observed in W . Note that if a state p has not been visited, $A_W(p, q)$ is estimated by $\frac{1}{m}$; otherwise, $A_W(p, q)$ is estimated by $\frac{n_{pq}}{n_p}$.

(3) *How do we know that the estimated transition probability matrix is accurate enough for the purpose of verifying fairness?* Formally, let A_W be the transition probability matrix estimated as above; and let A be the actual transition probability matrix. We would like the following to be satisfied.

$$P(\text{Div}(A, A_W) > \epsilon) \leq \delta \tag{4}$$

where $\epsilon > 0$ and $\delta > 0$ are constants representing *accuracy* and *confidence*; $Div(A, A_W)$ represents the divergence between A and A_W ; and P is the probability. Intuitively, the learned DTMC must be estimated such that the probability of the divergence between A_W and A greater than ϵ is no larger than the confidence level δ . In this work, we define the divergence based on the individual transition probability, i.e.,

$$P(\exists p \in S, \sum_{q \in S} |A(p, q) - A_W(p, q)| > \epsilon) \leq \delta \quad (5)$$

Intuitively, we would like to sample traces until the observed transition probabilities $A_W(p, q) = \frac{n_{pq}}{n_p}$ are close to the real transition probability $A(p, q)$ to a certain level for all $p, q \in S$. Theorem 1 in the recently published work [10] shows that if we sample enough samples such that for each $p \in S$, n_p satisfies

$$n_p \geq \frac{2}{\epsilon^2} \log\left(\frac{2}{\delta'}\right) \left[\frac{1}{4} - \left(\max_q \left| \frac{1}{2} - \frac{n_{pq}}{n_p} \right| - \frac{2}{3}\epsilon \right)^2 \right] \quad (6)$$

where $\delta' = \frac{\delta}{m}$, we can guarantee the learned DTMC is sound with respect to N in terms of probabilistic reachability analysis. Formally, let $H(n) = \frac{2}{\epsilon^2} \log\left(\frac{2}{\delta'}\right) \left[\frac{1}{4} - \left(\max_q \left| \frac{1}{2} - \frac{n_{pq}}{n_p} \right| - \frac{2}{3}\epsilon \right)^2 \right]$,

Theorem 1. *Let (S, I, A_W) be a DTMC where A_W is the transition probability matrix learned using frequency estimation based on n traces W . For $0 < \epsilon < 1$ and $0 < \delta < 1$, if for all $p \in S$, $n_p \geq H(n)$, we have for any CTL property ψ ,*

$$P(|\gamma(A, \psi) - \gamma(A_W, \psi)| > \epsilon) \leq \delta \quad (7)$$

where $\gamma(A_W, \psi)$ is the probability of A_W satisfying ψ .

Appendix A.3 in [50] provides the proof. Intuitively, the theorem provides a bound on the number of traces that we must sample in order to guarantee that the learned DTMC is PAC-correct with respect to any CTL property, which provides a way of verifying fairness as we show in Sect. 3.2.

We now go through Algorithm 2 in detail. The loop from line 3 to 8 keeps sampling inputs and obtains traces. Note that we use the uniform sampling by default and would sample according to the actual distribution if it is provided. Next, we update A_W as explained above at line 6. Then we check if more samples are needed by monitoring if a sufficient number of traces has been sampled according to Theorem 1. If it is the case, we output the DTMC as the result. Otherwise, we repeat the steps to generate new samples and update the model.

Example 2. In our running example, for simplicity assume that we select *gender* (as the protected feature) and the prediction outcome to be included in S and the number of clusters is set to 2 for both layers. Naturally, the two clusters identified for the protected feature are male and female (written as ‘ M ’ and ‘ F ’) and the two clusters determined for the outcome are ‘ $\leq 50K'$ ’ and ‘ $> 50K'$ ’.

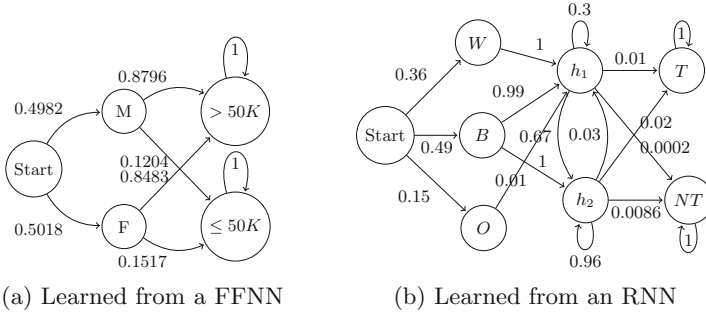


Fig. 1. Sample learned DTMCs

A sample trace is $w = \langle \text{Start}, 'M', '> 50K' \rangle$, where *Start* is a dummy state where all traces start. Assume that we set accuracy $\epsilon = 0.005$ and confidence level $\delta = 0.05$. Applying Algorithm 2, 2.85K traces are generated to learn the transition matrix A_W . The learned DTMC D is shown in Fig. 1a.

3.2 Step 2: Probabilistic Model Checking

In this step, we verify N against the fairness property ϕ based on the learned D . Note that D is PAC-correct only with respect to CTL properties. Thus it is infeasible to directly verify ϕ (which is not a CTL property). Our remedy is to compute $P(Y = l \mid F = f_i)$ and $P(Y = l \mid F = f_j)$ separately and then verify ϕ based on the results. Because we demand there is always a state in S representing $F = f_i$ and a state representing $Y = l$, the problem of computing $P(Y = l \mid F = f_i)$ can be reduced to a probabilistic reachability checking problem ψ , i.e., the probability of reaching the state representing $Y = l$ from the state representing $F = f_i$. This can be solved using probabilistic model checking techniques. Probabilistic model checking [39] of DTMC is a formal verification method for analyzing DTMC against formally-specified quantitative properties (e.g., PCTL). Probabilistic model checking algorithms are often based on a combination of graph-based algorithms and numerical techniques. For straightforward properties such as computing the probability that a **U** (Until), **F** (Finally) or **G** (*Globally*) path formula is satisfied, the problem can be reduced to solving a system of linear equations [39]. We refer to the readers to [39] for a complete and systematic formulate of the algorithm for probabilistic model checking.

Example 3. Figure 1b shows a DTMC learned from a recurrent neural network trained on Jigsaw Comments dataset (refer to details on the dataset and network in Sect. 4.1). The protected features is *race*. For illustration purpose, let us consider three different values for *race*, i.e., White (*W*), Black (*B*) and Others (*O*). For the hidden layer cells, we consider LSTM cell 1 only and cluster its values into two groups, represented as two states h_1 and h_2 . The output has two categories, i.e., Toxic (*T*) and Non-Toxic (*NT*). The transition probabilities are

shown in the figure. Note that the DTMC is cyclic due to the recurrent hidden LSTM cells in the network. We obtained $P(Y = 'T' | F = 'W')$ by probabilistic model checking as discussed above. The resultant probability is 0.0263. Similarly, $P(Y = 'T' | F = 'B')$ and $P(Y = 'T' | F = 'O')$ are 0.0362 and 0.0112 respectively.

Next we verify the fairness property ϕ based on the result of probabilistic model checking. First, the following is immediate based on Theorem 1.

Proposition 1. *Let $D = (S, I, A_W)$ be a DTMC learned using Algorithm 2. Let $P(Y = l | F = f_i)$ be the probability computed based on probabilistic model checking D and $P_t(Y = l | F = f_i)$ is the actual probability in N . We have*

$$P(|P(Y = l | F = f_i) - P_t(Y = l | F = f_i)| > \epsilon) \leq \delta$$

Theorem 2. *Let X be an estimation of a probability X_t such that $P(|X - X_t| > \epsilon) \leq \delta$. Let Z be an estimation of a probability Z_t such that $P(|Z - Z_t| > \epsilon) \leq \delta$. We have $P(|(X - Z) - (X_t - Z_t)| > 2\epsilon) \leq 2\delta - \delta^2$.*

Appendix A.4 in [50] provides the proof. Hence, given an expected accuracy $\mu\epsilon$ and a confidence level $\mu\delta$ on fairness property ϕ , we can derive ϵ and δ to be used in Algorithm 2 as: $\epsilon = \frac{\mu\epsilon}{2}$ and $\delta = 1 - \sqrt{1 - \mu\delta}$. We compute the probability of $P(Y = l | F = f_i)$ and $P(Y = l | F = f_j)$ based on the learned D (i.e., line 3 of Algorithm 1). Next, we compare $|P(Y = l | F = f_i) - P(Y = l | F = f_j)|$ with ξ . If the difference is no larger than ξ , fairness is verified. The following establishes the correctness of Algorithm 1.

Theorem 3. *Algorithm 1 is PAC-correct with accuracy $\mu\epsilon$ and confidence $\mu\delta$, if Algorithm 2 is used to learn the DTMC D .*

Appendix A.5 in [50] provides the proof.

The overall time complexity of model learning and probabilistic model checking is linear in the number of traces sampled, i.e., $\mathbf{O}(n)$ where n is the total number of traces sampled. Here n is determined by $H(n)$ as well as the probability distribution of the states. Contribution of $H(n)$ can be determined as $\mathbf{O}(\frac{\log m}{\mu\epsilon^2 \log \mu\delta})$ based on Eq. 6, where m is the total number of states. In the first case, for a model with only input features and output predictions as states, the probability of reaching each input states are statistically equal if we apply uniform sampling to generate IID input vectors. In this scenario the overall time complexity is $\mathbf{O}(\frac{m \log m}{\mu\epsilon^2 \log \mu\delta})$. In the second case, for a model with states representing the status of hidden layer neurons, we need to consider the probability for each hidden neuron states when the sampled inputs are fed into the network N . In the best case, the probabilities are equal, we denote m' as the maximum number of states in one layer among all layers included, the complexity is then $\mathbf{O}(\frac{m' \log m}{\mu\epsilon^2 \log \mu\delta})$. In the worst case, certain neuron is never activated (or certain predefined state is never reached) no matter what the input is. Since the probability distribution among the hidden states are highly network-dependent, we are not able to estimate the average performance.

Example 4. In our running example, with the learned A_W of D as shown in Fig. 1a, the probabilities as $P(Y = 1 | F = 'F') = 0.8483$ and $P(Y = 1 | F = 'M') = 0.8796$. Hence, $|P(Y = 1 | F = 'F') - P(Y = 1 | F = 'M')| = 0.0313$. Next, we compare the probability difference against the user-provided fairness criteria ξ . If $\xi = 0.1$, N satisfies fairness property. If $\xi = 0.02$, N fails fairness. Note that such a strict criteria is not practical and is used for illustration purpose only.

3.3 Step 3: Sensitivity Analysis

In the case that the verification result shows ϕ is satisfied, our approach outputs D and terminates successfully. We remark that in such a case D can be regarded as the evidence for the verification result as well as a human-interpretable explanation on why fairness is satisfied. In the case that ϕ is not satisfied, a natural question is: how do we improve the neural network for fairness? Existing approaches have proposed methods for improving fairness such as by training without the protected features [55] (i.e., a form of pre-processing) or training with fairness as an objective [16] (i.e., a form of in-processing). In the following, we show that a post-processing method can be supported based on the learned DTMC. That is, we can identify the neurons which are responsible for violating the fairness based on the learned DTMC and “mutate” the neural network slightly, e.g., adjusting its weights, to achieve fairness.

We start with a sensitivity analysis to understand the impact of each probabilistic distribution (e.g., of the non-protected features or hidden neurons) on the prediction outcome. Let F be the set of discrete states representing different protected feature values. Let I represent a non-protected feature or an internal neuron. We denote I_i as a particular state in the DTMC which represents certain group of values of the feature or neuron. Let l represent the prediction result that we are interested in. The sensitivity of I (with respect to the outcome l) is defined as follows.

$$\text{sensitivity}(I) = \sum_i \text{reach}(S_0, I_i) * \text{reach}(I_i, l) * \max_{\{f, g\} \subseteq F} (\text{reach}(f, I_i) - \text{reach}(g, I_i))$$

where $\text{reach}(s, s')$ for any state s and s' represents the probability of reaching s' from s . Intuitively, the sensitivity of I is the summation of the ‘sensitivity’ of every state I_i , which is calculated as $\max_{f, g} (\text{reach}(f, I_i) - \text{reach}(g, I_i))$, i.e., the maximum probability difference of reaching I_i from all possible protected feature states. The result is then multiplied with the probability of reaching I_i from start state S_0 and the probability of reaching l from I_i . Our approach analyzes all non-protected features and hidden neurons and identify the most sensitive features or neurons for improving fairness in step 4.

Example 5. In our running example, based on the learned DTMC D shown in Fig. 1a, we perform sensitivity analysis as discussed above. We observe that feature 9 (i.e., representing ‘capital gain’) is the most sensitive, i.e., it has the most contribution to the model unfairness. More importantly, it can be observed that the sensitivities of the neurons vary significantly, which is a good news as

it suggests that for this model, optimizing the weights of a few neurons may be sufficient for achieving fairness. Figure 3 in Appendix A.6 in [50] shows the sensitivity analysis scatter plot.

3.4 Step 4: Improving Fairness

In this step, we demonstrate one way of improving neural network fairness based on our analysis result, i.e., by adjusting weight parameters of the neurons identified in step 3. The idea is to search for a small adjustment through optimization techniques such that the fairness property is satisfied. In particular, we adopt the Particle Swarm Optimization (PSO) algorithm [37], which simulates intelligent collective behavior of animals such as flocks of birds and schools of fish. In PSO multiple particles are placed in the search space and the optimization target is to find the best location, where the fitness function is used to determine the best location. We omit the details of PSO here due to space limitation and present it in Appendix A.7 in [50].

In our approach, the weights of the most sensitive neurons are the subject for optimization and thus are represented by the location of the particles in the PSO. The initial location of each particle is set to the original weights and the initial velocity is set to zero. The fitness function is defined as follows.

$$fitness = Prob_{diff} + \alpha(1 - accuracy) \quad (8)$$

where $Prob_{diff}$ represents the maximum probability difference of getting a desired outcome among all different values of the sensitive feature; $accuracy$ is the accuracy of repaired network on the training dataset and constant parameter $\alpha \in (0, 1)$ determines the importance of the accuracy (relative to the fairness). Intuitively, the objective is to satisfy fairness and not to sacrifice accuracy too much. We set the bounds of weight adjustment to $(0, 2)$, i.e., 0 to 2 times of the original weight. The maximum number of iteration is set to 100. To further reduce the searching time, we stop the search as soon as the fairness property is satisfied or we fail to find a better location in the last 10 consecutive iterations.

Example 6. In our running example, we optimize the weight of ten most sensitive neurons using PSO for better fairness. The search stops at the 13rd iteration as no better location is found in the last 10 consecutive iterations. The resultant probability difference among the protected features dropped from 0.0313 to 0.007, whereas the model accuracy dropped from 0.8818 to 0.8606.

4 Implementation and Evaluation

Our approach has been implemented on top of SOCRATES [45], which is a framework for experimenting neural network analysis techniques. We conducted our experiments on a machine with 1 Dual-Core Intel Core i5 2.9GHz CPU and 8GB system memory.

Table 1. Fairness verification results

Dataset	Feature	# States	# Traces	Max Prob. Diff.	Result	Time
Census	Race	8	12500	0.0588	PASS	4.13 s
Census	Age	12	23500	0.0498	PASS	6.31 s
Census	Gender	5	2850	0.0313	PASS	0.98 s
Credit	Age	11	22750	0.1683	Fail	6.72 s
Credit	Gender	5	2850	0.0274	PASS	1.01 s
Bank	Age	12	27200	0.0156	PASS	6.33 s
Jigsaw	Religion	10	35250	0.0756	PASS	29.6 m
Jigsaw	Race	7	30550	0.0007	PASS	27.3 m

4.1 Experiment Setup

In the following, we evaluate our method in order to answer multiple research questions (RQs) based on multiple neural networks trained on 4 datasets adopted from existing research [4, 28, 54, 58], i.e., in addition to the *Census Income* [18] dataset as introduced in Example 1, we have the following three datasets. First is the *German Credit* [19] dataset consisting of 1k instances containing 20 features and is used to assess an individual’s credit. Age and gender are the two protected features. The labels are whether an individual’s credit is good or not. Second is the *Bank Marketing* [17] dataset consisting of 45k instances. There are 17 features, among which age is the protected feature. The labels are whether the client will subscribe a term deposit. Third is Jigsaw Comment [1] dataset. It consists of 313k text comments with average length of 80 words classified into toxic and non-toxic. The protected features analysed are race and religion.

Following existing approaches [4, 28, 54, 58], we train three 6-layer feed-forward neural networks (FFNN) on the first three dataset (with accuracy 0.88, 1 and 0.92 respectively) and train one recurrent neural network, i.e., 8-cell Long Short Term Memory (LSTM), for the last dataset (with accuracy 0.92) and analyze their fairness against the corresponding protected attributes. For the LSTM model, we adopt the state-of-the-art embedding tool GloVe [44]. We use the 50-dimension word vectors pre-trained on Wikipedia 2014 and Gigaword 5 dataset.

Recall that we need to sample inputs to learn a DTMC. In the case of first three datasets, inputs are sampled by generating randomly values within the range of each feature (in IID manner assuming a uniform distribution). In the case of the Jigsaw dataset, we cannot randomly generate and replace words as the resultant sentence is likely invalid. Inspired by the work in [7, 33, 41], our approach is to replace a randomly selected word with a randomly selected synonym (generated by Gensim [46]).

4.2 Research Questions and Answers

RQ1: Is our Approach Able to Verify Fairness? We systematically apply our method to the above-mentioned neural networks with respect to each protected

feature. Our experiments are configured with accuracy $\mu\epsilon = 0.01$, confidence level $\mu\delta = 0.1$ (i.e., $\epsilon = 0.005$, $\delta = 0.05$) and fairness criteria $\xi = 10\%$ (which is a commonly adopted threshold [6]). Furthermore, in this experiment, the states in the DTMC S are set to only include those representing the protected feature and different predictions. Table 1 summarizes the results. We successfully verify or falsify all models. Out of eight cases, the model trained on the German Credit dataset fails fairness with respect to the feature *age* (i.e., the maximum probability difference among different age groups is 0.1683 which is greater than $\xi = 10\%$). Furthermore, the model trained on the Jigsaw dataset shows some fairness concern with respect to the feature *religion* (although the probability difference is still within the threshold). This result shows that fairness violation could be a real concern.

RQ2: How Efficient is our Approach? We answer the question using two measurements. The first measurement is the execution time. The results are shown in the last column in Table 1. For the six cases of FFNN, the average time taken to verify a network is around 4.25s, with a maximum of 6.72s for the model trained on German Credit on feature *age* and a minimum of 0.98s for the model trained on the Census Income dataset on feature *gender*. For the two cases of LSTM networks, the average time taken is 30 min. Compared with FFNN, verifying an LSTM requires much more time. This is due to three reasons. Firstly, as mentioned in Sect. 4.1, sampling texts requires searching for synonyms. This is non-trivial due to the large size of the dictionary. Secondly, during sampling, we randomly select instances from the training set and apply perturbation to them in order to generate new samples. However, most of the instances in the Jigsaw training set does not contain the sensitive word. This leads to an increased number of traces needed to learn a DTMC. Thirdly, the LSTM model takes much more time to make a prediction than that by FFNN in general. It is also observed that for all the cases, the execution time is proportional to the number of traces used in DTMC model learning (as discussed in our time complexity analysis). The other measurement is the number of traces that we are required to sample using Algorithm 2. For each model and protected feature, the number of traces generated in Algorithm 2 depends on number of categorical values defined for this protected feature and the number of predictions. That is, more

Table 2. Fairness improvement

Dataset	Max probability difference	Accuracy
German	0.1683 \mapsto 0.1125	1.0 \mapsto 0.9450
Census	0.0588 \mapsto 0.0225	0.8818 \mapsto 0.8645
Jigsaw	0.0756 \mapsto 0.0590	0.9166 \mapsto 0.9100

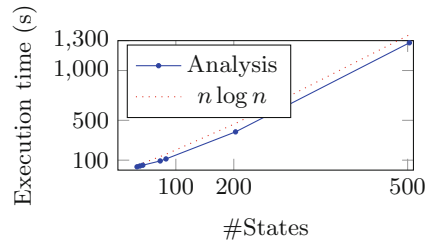


Fig. 2. Execution times vs number of states

categories and predictions result in more states in the learn a DTMC model, which subsequently lead to more traces required. Furthermore, the number of traces required also depends on the probabilistic distribution from each state in the model. As described in Algorithm 2, the minimum number of traces transiting from each state must be greater than $H(n)$. This is evidenced by results shown in Table 1, where the number of traces vary significantly between models or protected features, ranging from 2 K to 35 K. Although the number of traces is expected to increase for more complicated models, we believe that this is not a limiting factor since the sampling of the traces can be easily paralleled.

We further conduct an experiment to monitor the execution time required for the same neural network model with a different numbers of states in the learned DTMC. We keep other parameters (i.e., $\mu\epsilon$, $\mu\delta$ and ϕ) the same. Note that hidden neurons are not selected as states to reduce the impact of the state distribution. We show one representative result (based on the mode trained on the Census Income dataset with attribute *race* as the protected feature) in Fig. 2. As we can see the total execution time is bounded by $n \log n$ which tally with our time complexity analysis in Sect. 3.

RQ3: Is our Approach able to Improve Fairness and is the Sensitivity Analysis Useful? The question asks whether the sensitivity analysis results based on the learned DTMC can be used to improve fairness. To answer this question, we systematically perform sensitivity analysis (on both the input features and the hidden neurons) and optimize the weights of the neurons which are sensitive to fairness. We focus on three cases, i.e., the FFNN model trained on the German Credit model w.r.t *age* and on the Census Income model w.r.t *race* and the LSTM model trained on the Jigsaw comments w.r.t *religion*, as the maximum probability difference for these three cases (as shown in Table 1) is concerning (i.e., $> 5\%$). For the former two, we optimize the weights of the top-10 sensitive neurons (including the first layer neurons representing other features). For the LSTM model, we optimize top-3 sensitive cells (due to the small number of cells). Table 2 shows the fairness improvement as well as the drop in accuracy. It can be observed that in all three cases we are able to improve the fairness whilst maintaining the accuracy at a high-level. Note that the parameter α is set to 0.1 in these experiments and it can be used to achieve better fairness or accuracy depending the user requirement.

RQ4: How does our Approach Compare with Existing Alternatives? The most relevant tools that we identify are FairSquare [6] and VeriFair [9]. FairSquare and VeriFair use numerical integration to verify fairness properties of machine learning models including neural networks. FairSquare relies on constraint solving techniques and thus it is difficult to scale to large neural networks. VeriFair is based on adaptive concentration inequalities. We evaluate our approach against these two tools on all eight models. For FairSquare and VeriFair, we follow the setting of independent feature distribution and check for demographic parity [9]. For both tools, we set $c = 0.15$ as suggested and keep other parameters as default. As both FairSquare and VeriFair are designed to compare two groups of sub-populations, for those protected features that have more than two

Table 3. Comparison with FairSquare, VeriFair and Statistical Model Checking (SMC)

Dataset	Prot. Feat.	FairSquare		VeriFair		SMC		Ours	
		Result	Time	Result	Time	Result	Time	Result	Time
Census	Race	–	T.O.	Pass	2.33 s	Pass	0.81 s	Pass	0.93 s
Census	Age	–	T.O.	Pass	37.14 s	Pass	0.74 s	Pass	0.81 s
Census	Gender	–	T.O.	Pass	2.19 s	Pass	0.69 s	Pass	0.89 s
Credit	Age	–	T.O.	Pass	39.29 s	Pass	0.78 s	Pass	0.90 s
Credit	Gender	–	T.O.	Pass	8.23 s	Pass	0.70 s	Pass	0.82 s
Bank	Age	–	T.O.	Pass	245.34 s	Pass	0.66 s	Pass	0.97 s
Jigsaw	Religion	–	–	–	–	Pass	14.9 m	Pass	15.3 m
Jigsaw	Race	–	–	–	–	Pass	14.2 m	Pass	14.5 m

categories, we perform binning to form two groups. For the six FFNN models, we set timeout value to be 900s following the setting in VeriFair. As shown in Table 3, FairSquare is not able to scale for large neural network and for all FFNN models it fails to verify or falsify the model in time. Both VeriFair and our approach successfully verified all six FFNN models. But our approach completes the verification within 1s for all models while VeriFair takes 62 times more execution time than our approach on average. For the RNN models trained on Jigsaw dataset, neither FairSquare nor VeriFair is able to analyze them. FairSquare supports only loop-free models and, hence, it cannot handle RNN models. Although VeriFair is able to handle RNN networks in general, it does not support text classifiers. To evaluate the overhead introduced by constructing DTMCs in our approach, we further compare our approach with statistical model checking (SMC). In this experiment, we generate the same amount of samples as required by our approach for each scenario and directly perform statistical model checking. The result is shown in 3 column *SMC*. Compared with our approach, it is observed that most of the time were taken by sampling and inference and the overhead of building DTMCs is around 2.4% on average. Hence, compared with existing solutions, our approach is more efficient than FairSquare and VeriFair and can support RNN-based text classifiers. Furthermore, the overhead of constructing DTMCs is negligible.

5 Related Work

Neural network verification. There have been multiple approaches proposed to verify the robustness of neural networks utilizing various techniques, i.e., abstraction [23, 29, 48], SMT solver [32, 34, 35], MILP and LP [22, 52], symbolic execution [57] and many others [15, 20, 31, 38]. Besides [6] and [9] that we addressed in RQ4, [28] and [43] studied fairness property of text classifiers. Unlike ours, they focus on text classifiers only and their performance on RNN is unknown.

Fairness testing and improvement. There have been an impressive line of methods proposed recently on machine learning model fairness testing and

improvement. THEMIS [8,27], AEQUITAS [54], Symbolic Generation (SG) [4] and ADF [58], are proposed to generate discriminatory instances for fairness testing. There are also existing proposals on fairness training [3,12,14,16,28,36]. Our work instead focuses on post-processing where a trained model is repaired based on sensitivity analysis results to improve fairness.

Machine learning model repair. There have been multiple approaches proposed to repair machine learning models based on various technologies, i.e., [49] leverages SMT solving, [30] is based on advances in verification methods, [5] is guided by input population and etc. Unlike these methods, our work focuses on fairness repair and supports FFNN and RNN by design.

6 Conclusion

In this work, we proposed an approach to formally verify neural networks against fairness properties. Our work relies on an approach for learning DTMC from given neural network with PAC-correct guarantee. Our approach further performs sensitivity analysis on the neural network if it fails the fairness property and provides useful information on why the network is unfair. This result is then used as a guideline to adjust network parameters and achieve fairness. Comparing with existing methods evaluating neural network fairness, our approach has significantly better performance in terms of efficiency and effectiveness.

Acknowledgements. We thank anonymous reviewers for their constructive feedback. This research is partly supported by Singapore Ministry of Education (project MOET 32020 – 0004) and Guangdong Science and Technology Department (Grant Number 2018B010107004).

References

1. <https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge>
2. Draft ethics guidelines for trustworthy AI. Tech. rep., European Commission (2018)
3. Agarwal, A., Beygelzimer, A., Dudík, M., Langford, J., Wallach, H.M.: A reductions approach to fair classification. In: Dy, J.G., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Stockholm, Sweden, 10-15 July 2018. Proceedings of Machine Learning Research, vol. 80, pp. 60–69. PMLR (2018). <http://proceedings.mlr.press/v80/agarwal18a.html>
4. Agarwal, A., Lohia, P., Nagar, S., Dey, K., Saha, D.: Automated test generation to detect individual discrimination in AI models. CoRR (2018). <http://arxiv.org/abs/1809.03260>
5. Albarghouthi, A., D’Antoni, L., Drews, S.: Repairing decision-making programs under uncertainty. In: Majumdar, R., Kunčák, V. (eds.) Computer Aided Verification, CAV 2017. LNCS, vol. 10426, pp. 181–200. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63387-9_9
6. Albarghouthi, A., D’Antoni, L., Drews, S., Nori, A.V.: Fairsquare: probabilistic verification of program fairness. Proc. ACM Program. Lang. **1**(OOPSLA), 80:1–80:30 (2017). <https://doi.org/10.1145/3133904>

7. Alzantot, M., Sharma, Y., Elgohary, A., Ho, B., Srivastava, M.B., Chang, K.: Generating natural language adversarial examples. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP 2018), Brussels, Belgium, pp. 2890–2896 (2018). <https://doi.org/10.18653/v1/d18-1316>
8. Angell, R., Johnson, B., Brun, Y., Meliou, A.: Themis: automatically testing software for discrimination. In: Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/SIGSOFT FSE 2018), Lake Buena Vista, FL, USA, pp. 871–875 (2018). <https://doi.org/10.1145/3236024.3264590>
9. Bastani, O., Zhang, X., Solar-Lezama, A.: Probabilistic verification of fairness properties via concentration. PACMPL **3**(OOPSLA), 118:1–118:27 (2019). <https://doi.org/10.1145/3360544>
10. Bazille, H., Genest, B., Jegourel, C., Sun, J.: Global PAC bounds for learning discrete time Markov chains. In: Lahiri, S.K., Wang, C. (eds.) Computer Aided Verification, CAV 2020. LNCS, vol. 12225, pp. 304–326. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-53291-8_17
11. Ben-Ari, M., Pnueli, A., Manna, Z.: The temporal logic of branching time. Acta Informatica **20**, 207–226 (1983). <https://doi.org/10.1007/BF01257083>
12. Berk, R., et al.: A convex framework for fair regression. CoRR abs/1706.02409 (2017). <http://arxiv.org/abs/1706.02409>
13. Bojarski, M., et al.: End to end learning for self-driving cars. CoRR (2016). <http://arxiv.org/abs/1604.07316>
14. Bolukbasi, T., Chang, K., Zou, J.Y., Saligrama, V., Kalai, A.T.: Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In: Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016 (NeurIPS 2016), Barcelona, Spain, pp. 4349–4357 (2016). <http://papers.nips.cc/paper/6228-man-is-to-computer-programmer-as-woman-is-to-homemaker-debiasing-word-embeddings>
15. Bunel, R., Lu, J., Turkaslan, I., Torr, P.H.S., Kohli, P., Kumar, M.P.: Branch and bound for piecewise linear neural network verification. J. Mach. Learn. Res. **21**, 42:1–42:39 (2020). <http://jmlr.org/papers/v21/19-468.html>
16. Cava, W.L., Moore, J.: Genetic programming approaches to learning fair classifiers. In: Proceedings of the 2020 Genetic and Evolutionary Computation Conference (2020)
17. Dua, D., Graff, C.: Bank marketing dataset at UCI machine learning repository (2017). <https://archive.ics.uci.edu/ml/datasets/Bank+Marketing>
18. Dua, D., Graff, C.: Census income dataset at UCI machine learning repository (2017). <https://archive.ics.uci.edu/ml/datasets/adult>
19. Dua, D., Graff, C.: German credit dataset at UCI machine learning repository (2017). [https://archive.ics.uci.edu/ml/datasets/statlog+\(german+credit+data\)](https://archive.ics.uci.edu/ml/datasets/statlog+(german+credit+data))
20. Dvijotham, K.D., Stanforth, R., Goyal, S., Qin, C., De, S., Kohli, P.: Efficient neural network verification with exactness characterization. In: Globerson, A., Silva, R. (eds.) Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, 22–25 July 2019. Proceedings of Machine Learning Research, vol. 115, pp. 497–507. AUAI Press (2019). <http://proceedings.mlr.press/v115/dvijotham20a.html>
21. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.S.: Fairness through awareness. In: Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, pp. 214–226 (2012). <https://doi.org/10.1145/2090236.2090255>
22. Ehlers, R.: Formal verification of piece-wise linear feed-forward neural networks. CoRR abs/1705.01320 (2017). <http://arxiv.org/abs/1705.01320>

23. Elboher, Y.Y., Gottschlich, J., Katz, G.: An abstraction-based framework for neural network verification. In: Lahiri, S.K., Wang, C. (eds.) CAV 2020. LNCS, vol. 12224, pp. 43–65. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-53288-8_3
24. Feldman, M., Friedler, S.A., Moeller, J., Scheidegger, C., Venkatasubramanian, S.: Certifying and removing disparate impact. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, pp. 259–268 (2015). <https://doi.org/10.1145/2783258.2783311>
25. Fu, K., Cheng, D., Tu, Y., Zhang, L.: Credit card fraud detection using convolutional neural networks. In: Hirose, A., Ozawa, S., Doya, K., Ikeda, K., Lee, M., Liu, D. (eds.) Neural Information Processing, ICONIP 2016. LNCS, vol. 9949, pp. 483–490. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46675-0_53
26. Gale, W.: Good-turing smoothing without tears. *J. Quant. Linguist.* 217–37 (1995)
27. Galhotra, S., Brun, Y., Meliou, A.: Fairness testing: testing software for discrimination. In: Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2017), Paderborn, Germany, pp. 498–510 (2017). <https://doi.org/10.1145/3106237.3106277>
28. Garg, S., Perot, V., Limtiaco, N., Taly, A., Chi, E.H., Beutel, A.: Counterfactual fairness in text classification through robustness. In: Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (AIES 2019), Honolulu, HI, USA, pp. 219–226 (2019). <https://doi.org/10.1145/3306618.3317950>
29. Gehr, T., Mirman, M., Drachler-Cohen, D., Tsankov, P., Chaudhuri, S., Vechev, M.T.: AI2: safety and robustness certification of neural networks with abstract interpretation. In: 2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21–23 May 2018, San Francisco, California, USA, pp. 3–18. IEEE Computer Society (2018). <https://doi.org/10.1109/SP.2018.00058>
30. Goldberger, B., Katz, G., Adi, Y., Keshet, J.: Minimal modifications of deep neural networks using verification. In: Albert, E., Kovács, L. (eds.) LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, 22–27 May 2020. EPiC Series in Computing, vol. 73, pp. 260–278. EasyChair (2020). <https://easychair.org/publications/paper/CWhF>
31. Gross, D., Jansen, N., Pérez, G.A., Raaijmakers, S.: Robustness verification for classifier ensembles. In: Hung, D.V., Sokolsky, O. (eds.) Automated Technology for Verification and Analysis, ATVA 2020. LNCS, vol. 12302, pp. 271–287. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59152-6_15
32. Jacoby, Y., Barrett, C., Katz, G.: Verifying recurrent neural networks using invariant inference. In: Hung, D.V., Sokolsky, O. (eds.) Automated Technology for Verification and Analysis, ATVA 2020. LNCS, vol. 12302, pp. 57–74. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59152-6_3
33. Jia, R., Liang, P.: Adversarial examples for evaluating reading comprehension systems. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP 2017), Copenhagen, Denmark, pp. 2021–2031 (2017). <https://doi.org/10.18653/v1/d17-1215>
34. Katz, G., Barrett, C., Dill, D.L., Julian, K., Kochenderfer, M.J.: Reluplex: an efficient SMT solver for verifying deep neural networks. In: Majumdar, R., Kunčák, V. (eds.) Computer Aided Verification, CAV 2017. LNCS, vol. 10426, pp. 97–117. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63387-9_5

35. Katz, G., et al.: The marabou framework for verification and analysis of deep neural networks. In: Dillig, I., Tasiran, S. (eds.) *Computer Aided Verification, CAV 2019*. LNCS, vol. 11561, pp. 443–452. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-25540-4_26
36. Kearns, M.J., Neel, S., Roth, A., Wu, Z.S.: Preventing fairness gerrymandering: auditing and learning for subgroup fairness. In: Dy, J.G., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, 10–15 July 2018*. *Proceedings of Machine Learning Research*, vol. 80, pp. 2569–2577. PMLR (2018). <http://proceedings.mlr.press/v80/kearns18a.html>
37. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: *Proceedings of ICNN'95 - International Conference on Neural Networks*, vol. 4, pp. 1942–1948 (1995). <https://doi.org/10.1109/ICNN.1995.488968>
38. Ko, C., Lyu, Z., Weng, L., Daniel, L., Wong, N., Lin, D.: POPQORN: quantifying robustness of recurrent neural networks. In: Chaudhuri, K., Salakhutdinov, R. (eds.) *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9–15 June 2019, Long Beach, California, USA*. *Proceedings of Machine Learning Research*, vol. 97, pp. 3468–3477. PMLR (2019). <http://proceedings.mlr.press/v97/ko19a.html>
39. Kwiatkowska, M., Norman, G., Parker, D.: Advances and challenges of probabilistic model checking. In: *2010 48th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2010* (2010). <https://doi.org/10.1109/ALLERTON.2010.5707120>
40. Legay, A., Delahaye, B., Bensalem, S., et al.: Statistical model checking: an overview. In: Barringer, H. (ed.) *Runtime Verification*, pp. 122–135. Springer, Berlin Heidelberg, Berlin, Heidelberg (2010)
41. Li, J., Ji, S., Du, T., Li, B., Wang, T.: Textbugger: generating adversarial text against real-world applications. In: *26th Annual Network and Distributed System Security Symposium (NDSS 2019)*, San Diego, California, USA (2019). <https://www.ndss-symposium.org/ndss-paper/textbugger-generating-adversarial-text-against-real-world-applications/>
42. Lloyd, S.P.: Least squares quantization in PCM. *IEEE Trans. Inf. Theor.* **28**(2), 129–136 (1982). <https://doi.org/10.1109/TIT.1982.1056489>
43. Ma, P., Wang, S., Liu, J.: Metamorphic testing and certified mitigation of fairness violations in NLP models. In: Bessiere, C. (ed.) *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 458–465 (2020). <https://doi.org/10.24963/ijcai.2020/64>
44. Pennington, J., Socher, R., Manning, C.D.: Glove: global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, 25–29 October 2014, Doha, Qatar, pp. 1532–1543 (2014). <https://www.aclweb.org/anthology/D14-1162/>
45. Pham, L.H., Li, J., Sun, J.: SOCRATES: towards a unified platform for neural network verification. *CoRR abs/2007.11206* (2020). <https://arxiv.org/abs/2007.11206>
46. Řehůřek, R., Sojka, P.: Software framework for topic modelling with large Corpora. In: *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, Valletta, Malta, pp. 45–50 (2010). <http://is.muni.cz/publication/884893/en>

47. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: a unified embedding for face recognition and clustering. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2015), Boston, MA, USA, pp. 815–823 (2015). <https://doi.org/10.1109/CVPR.2015.7298682>
48. Singh, G., Gehr, T., Püschel, M., Vechev, M.T.: An abstract domain for certifying neural networks. *Proc. ACM Program. Lang.* **3**(POPL), 41:1–41:30 (2019). <https://doi.org/10.1145/3290354>
49. Sotoudeh, M., Thakur, A.: Correcting deep neural networks with small, generalizing patches. In: Workshop on Safety and Robustness in Decision Making (2019)
50. Sun, B., Sun, J., Dai, T., Zhang, L.: Probabilistic verification of neural networks against group fairness. *CoRR* abs/2107.08362 (2021). <https://arxiv.org/abs/2107.08362>
51. Thomas, P.S., da Silva, B.C., Barto, A.G., Giguere, S., Brun, Y., Brunskill, E.: Preventing undesirable behavior of intelligent machines. *Science* **366**(6468), 999–1004 (2019). <https://science.sciencemag.org/content/366/6468/999>
52. Tjeng, V., Xiao, K.Y., Tedrake, R.: Evaluating robustness of neural networks with mixed integer programming. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, 6–9 May 2019 (2019). <https://openreview.net/forum?id=HyGIiRqtm>
53. Tramèr, F., et al.: Fairtest: discovering unwarranted associations in data-driven applications. In: 2017 IEEE European Symposium on Security and Privacy (EuroS and P 2017), Paris, France, pp. 401–416 (2017). <https://doi.org/10.1109/EuroSP.2017.29>
54. Udeshi, S., Arora, P., Chattopadhyay, S.: Automated directed fairness testing. In: Huchard, M., Kästner, C., Fraser, G. (eds.) Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, 3–7 September 2018, pp. 98–108. ACM (2018). <https://doi.org/10.1145/3238147.3238165>
55. Veale, M., Binns, R.: Fairer machine learning in the real world: mitigating discrimination without collecting sensitive data. *Big Data Soc.* **4** (2017)
56. Vieira, S., Pinaya, W.H., Mechelli, A.: Using deep learning to investigate the neuroimaging correlates of psychiatric and neurological disorders: methods and applications. *Neurosci. Biobehav. Rev.* **74**, 58–75 (2017). <https://doi.org/10.1016/j.neubiorev.2017.01.002>
57. Wang, S., Pei, K., Whitehouse, J., Yang, J., Jana, S.: Formal security analysis of neural networks using symbolic intervals. In: Enck, W., Felt, A.P. (eds.) 27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, 15–17 August 2018, pp. 1599–1614. USENIX Association (2018). <https://www.usenix.org/conference/usenixsecurity18/presentation/wang-shiqi>
58. Zhang, P., et al.: White-box fairness testing through adversarial sampling. In: Proceedings of the 42th International Conference on Software Engineering (ICSE 2020), Seoul, South Korea (2020)