

FairyTED: A Fair Rating Predictor for TED Talk Data

Rupam Acharyya,¹ Shouman Das,¹ Ankani Chattoraj,¹ Md. Iftekhar Tanveer²

¹University of Rochester, ²Comcast Applied AI Research
racharyy@cs.rochester.edu, shouman.das@rochester.edu,
achattor@ur.rochester.edu, mdiftekhar_tanveer@comcast.com

Abstract

With the recent trend of applying machine learning in every aspect of human life, it is important to incorporate fairness into the core of the predictive algorithms. We address the problem of predicting the quality of public speeches while being fair with respect to sensitive attributes of the speakers, e.g. *gender* and *race*. We use the TED talks as an input repository of public speeches because it consists of speakers from a diverse community and has a wide outreach. Utilizing the theories of *Causal Models*, *Counterfactual Fairness* and state-of-the-art neural language models, we propose a mathematical framework for fair prediction of the public speaking quality. We employ grounded assumptions to construct a causal model capturing how different *attributes* affect public speaking quality. This causal model contributes in generating counterfactual data to train a *fair* predictive model. Our framework is general enough to utilize any assumption within the causal model. Experimental results show that while prediction accuracy is comparable to recent work on this dataset, our predictions are counterfactually fair with respect to a novel metric when compared to true data labels. The FairyTED setup not only allows organizers to make informed and diverse selection of speakers from the unobserved counterfactual possibilities but it also ensures that viewers and new users are not influenced by unfair and unbalanced ratings from arbitrary visitors to the ted.com website when deciding to view a talk.

1 Introduction

In recent times, artificial intelligence is being used in consequential decision making. Governments make use of it in criminal justice system to predict recidivism (Brennan, Dieterich, and Ehret 2009; Tollenaar and Van der Heijden 2013) which affects the decision about bail, sentencing and parole. Various firms are also using machine learning algorithms to examine and filter resumes of job applicants (Nguyen and Gatica-Perez 2016; Chen et al. 2017a; Naim et al. 2016) which is crucial for the growth of a company. Machine learning algorithms are also being used to evaluate human’s social skills such as presentation performance (Chen et al. 2017b; Tanveer, Liu, and Hoque 2015),

essay grading (Alikaniotis, Yannakoudakis, and Rei 2016; Taghipour and Ng 2016) etc.

To solve such decision making problems, machine learning algorithms are trained on massive datasets that are usually collected in the wild. Due to difficulties in the manual curation or adjustment over large dataset, it is likely that the data capture unwanted bias towards the underrepresented group based on race, gender or ethnicity. Such bias results in unfair decision making systems, leading to unwanted and often catastrophic consequences to human life and society. For example, the recognition rates of pedestrians in autonomous vehicles are reported to be not equally accurate for all groups of people (Wilson, Hoffman, and Morgenstern 2019). Matthew et al. (Kay, Matuszek, and Munson 2015) showed that societal bias gets reflected in the machine learning algorithms through biased dataset and causes representational harm for occupations. Face recognition has been found to be not as effective for people with different skin tones. Dark-skinned females have 43 times higher detection error than light-skinned males (Buolamwini and Gebru 2018).

In this work, we propose a predictive framework that tackles the issue of designing a fair prediction system from biased data. As an application scenario, we choose the problem of fair rating prediction in the TED talks. TED talks cover a wide variety of topics and influence audience by educating and inspiring them. In addition, it consists of speakers from a diverse community with imbalances in the age, gender and ethnic attributes. The ratings are provided by spontaneous visitors to the TED talk website. A machine learning algorithm trained solely from the audience ratings will have a possibility of the predicted rating being biased by sensitive attributes of the speakers.

It is a challenging problem because human behavior is driven by numerous factors and hence have huge variability. It is difficult to know the way these factors interact among each other. In addition, uncovering the true interaction model may not be feasible and often expensive. Even though the sharing platforms such as YouTube, Massive Open Online Courses (MOOC), or ted.com make it possible to collect a large amount of observational data, these platforms do not correct for bias and unfair ratings.

In this work, we utilize *causal models* (Pearl and others 2009) to define possible dependencies between attributes of the data. We then address the issue of not knowing the true interaction model by averaging outputs of predictors across several possible causes. Further using these causal models we generate *counterfactual samples* of the sensitive attributes. These counterfactual samples are the key components in our fair prediction framework (adapted from Kusner et al. (2017)) and help reducing bias in the ratings with respect to sensitive attributes. Finally, we introduce a novel metric to quantify the degree of fairness employed by our FairyTED pipeline. To the best of our knowledge, FairyTED is the first fair prediction pipeline for public speaking dataset and can be applied to any dataset of similar grounds. Apart from the theoretical contribution, our work also has practical implications in helping both the viewers and the organizers make informed and unbiased choices for selection of talks and speakers.

2 Related Works

There has been a rising interest in developing fair algorithms focused to mitigate the bias arising from discriminatory preferences of attributes such as gender, ethnicity, race, etc. These can cause bias across various domain, from college admissions process to criminal justice (Bickel, Hammel, and O’Connell 1975; Brennan, Dieterich, and Ehret 2009). Training a machine learning algorithm with an objective of getting higher prediction accuracy can be sometimes unfair towards underrepresented groups in the dataset. For example, it has been shown in (Bolukbasi et al. 2016) that the geometry of word embedding trained with traditional machine learning algorithms reflect gender stereotypes present in our society. Since machine learning models are used to take important and sensitive decisions including credit score prediction, loan applications assessment or predicting crime scenes, a careful approach should be designed to make traditional predictive model fair. Baeza-Yates emphasized the importance of increasing awareness about fairness in web based system. Calders and Verwer (2010) proposed methods for designing discrimination-free bayesian classifier. Dwork et al. (2012) formulated fairness as an optimization problem and made use of a task specific similarity metric which describes the similarity of two individuals for the classification task. Grgic-Hlaca et al. (2016) defined the notion of process fairness by focusing on the process of decision making rather than outcome of the classifier. Schumann et al. (2019) proposed a framework with theoretical guarantees to transfer fairness in machine learning across various domains. Other related research has been done where the main focus is to quantify the unfairness in a machine learning algorithms and create a model for a certain dataset. Readers are referred to (Kamiran and Calders 2009; Kamishima, Akaho, and Sakuma 2011; Joseph et al. 2016; Garg et al. 2019). For a recent complete survey see (Mehrabi et al. 2019). The notion of fairness in a prediction algorithm has been defined in various ways based on researcher’s assumption of fairness (Zliobaite 2015; Zafar et al. 2017). We follow the causal approach first introduced by Kusner et al. (2017) to address the notion of fairness in a machine

learning model. For the causal framework, we have adopted the definition of (Pearl and others 2009).

3 Preliminaries

3.1 Causal Model Definition

Following general convention we define causal model as a Directed Acyclic Graph (DAG) with a set of nodes (N) and edges ($E \subseteq N \times N$). Let $Pa_i = \{n_j | (n_j, n_i) \in E\}$ denote the set of parents of node n_i . Adapting the conventions used in Pearl and others (2009), Kusner et al. (2017), we then define the main characteristics of the causal DAGs. Each causal DAG consists of the triple $(\mathcal{U}, \mathcal{V}, \mathcal{F})$ where,

- \mathcal{U} denotes the set of unobserved variables in the outside world that influence observed variables of the causal models.
- \mathcal{V} denotes the set of observed variables consisting of three mutually exclusive sets, 1) set of sensitive attributes S , 2) set of data attributes X and 3) label Y ; i.e, $\mathcal{V} = S \cup X \cup Y$.
- $\mathcal{F} = \{F_1, F_2, \dots, F_n\}$ is a set of functions that define the relationship between $v_i \in \mathcal{V}$ and Pa_i for all i . In other words, $v_i = F_i(Pa_i) + \eta_i$, where η_i is a random variable drawn from a distribution and $Pa_i \subseteq \mathcal{V} \cup \mathcal{U}$.

Defining the set of functions \mathcal{F} is crucial for generating counterfactual samples of S and performing related computations (for details see Kusner et al. (2017) and Pearl and others (2009)). For intuitive examples explaining the variables $(\mathcal{U}, \mathcal{V}, \mathcal{F})$ in relation to causal models, see Section 6 and Fig. 3.

3.2 Counterfactual sample generation

To create an augmented dataset including actual observations and counterfactual samples of S we take the following steps (see chapter 4 of Pearl and others):

- We assume a prior distribution over \mathcal{U} and infer its posterior distribution given \mathcal{V} .
- Intervene on sensitive attributes $S \subseteq \mathcal{V}$ to generate counterfactual samples of S . The counterfactual samples of S are then augmented with actual observations to create augmented dataset (See Section 6 for details).

3.3 Counterfactual fairness

We adapt the definition of counterfactual fairness (Russell et al. 2017) and use $Y_{S \rightarrow s'}$ to denote the label of the counterfactual sample of S . A predictor \hat{Y} of Y is said to be counterfactually fair given the observed data attributes X and sensitive attributes S if $P(\hat{Y}_{S \leftarrow s} = y | X = x, S = s) = P(\hat{Y}_{S \leftarrow s'} = y | X = x, S = s)$ for all $s' \neq s$ and all y . Intuitively, this equation ensures that the prediction probability remains unaffected by interventions on sensitive attributes S when all other attributes are same. For example, if we observe that talks given by white male speakers are rated to be fascinating with a probability of 0.6 then counterfactually assigning the same talk content and other attributes to white females, say, should not change the probability from 0.6.

Table 1: TED talk Dataset Properties: Information about the TED talk videos that are used in the causal DAGs

Property	Quantity
Total number of talks	2,383
Total number of views	4206,164,936
Total length of all talks	564.63 Hours
Total number of ratings	5,954,233

4 Data

The data analyzed in our study was collected from TED talk website (ted.com). We crawled the website to obtain data from 2400 videos published between 2006 and 2017, covering a wide range of topics such as cultural, social and scientific issues. This not only highlights the vast appeal and diversity of TED talks but also exhibits the importance of fair rating predictions. We have removed 17 talks from our dataset as those were not held at a public speaking set up. The preliminary data contains details about the total number of *views* (V), the *transcripts* (T) used by the speaker, *ratings* (Y) of the videos given by the viewers, etc. The rating Y for each video consists of 14 labels such as *beautiful*, *ingenious* and *confusing*. Summary of the dataset is given in Table 1.

4.1 Data Annotations

We use Amazon mechanical turk to collect data on the protected attributes S (*race* and *gender*). Each video was annotated by 3 turkers and we verified the inter-rater reliability using Krippendorff’s alpha (Krippendorff 2011) which gives an average agreement of 93%. The remaining data was manually investigated and annotated.

4.2 Data Preprocessing

- We obtain embedded transcript ($T \in \mathbb{R}^d$) using the doc2vec implementation of Gensim package (Le and Mikolov 2014) and use $d = 200$ for all reported results.
- The original view count ($V_{old} \in \mathbb{Z}$) in the data ranges over large values compared to other attributes. We use the min-max technique to normalize V_{old} and obtain $V \in \mathbb{R}$.

$$V = \frac{V_{old} - \min\{V_{old}\}_{talks}}{\max\{V_{old}\}_{talks} - \min\{V_{old}\}_{talks}}$$

We assume that, how long a video has been online gets inherently captured by the total views and does not need to be explicitly modeled.

- We denote each original rating as $Y_{(old)} = (y_{1(old)}, \dots, y_{14(old)}) \in \mathbb{Z}^{14}$, where $y_{i(old)}$ is the count of i^{th} label from viewers. $Y_{(old)}$ is scaled w.r.t corresponding total ratings to acquire $Y \in \mathbb{R}^{14}$ as,

$$y_i = \frac{y_{i(old)}}{\sum_{j=1}^{14} y_{j(old)}}$$

- We binarize each rating label y_j by thresholding w.r.t median m_j ($median\{y_j\}_{talks}$). For each j the label y_j then becomes 0 or 1, where 1 indicates the $y_j > m_j$. We train

Table 2: Pre-processed Dataset Attributes: Final attributes from the TED talk dataset that are used to train the FairyTED classifier

Sensitive attributes	S , race and gender
Data attributes	X , transcript and view count
Label	Y , rating

our classifier to predict these 14 binarized rating labels. The attributes of our final dataset are shown in Table 2.

5 Observation in Data

We used an open source tool-kit AIF360 (Bellamy et al. 2018) to examine existing *bias* or *unfairness* in our preprocessed dataset w.r.t. S (race and gender of the speaker). We calculated the *statistical parity difference* (SPD) and *disparate impact* (DI) (Biddle 2006) for each of the 14 binarized rating labels as,

$$\begin{aligned} \text{SPD} &= \mathbb{P}(y_i = 1 | S \in \text{Grp } 1) - \mathbb{P}(y_i = 1 | S \in \text{Grp } 2) \\ \text{DI} &= \mathbb{P}(y_i = 1 | S \in \text{Grp } 1) / \mathbb{P}(y_i = 1 | S \in \text{Grp } 2) \end{aligned}$$

Using these metrics we calculate the marginal probability of y_i for each i across various groups and observe many significant differences. Fig. 1 shows some examples where we compare *male* speakers with speakers of *other genders*. The difference between blue and orange bars are noticeable for rating labels marked with red blocks. We observe that talks from male speakers are rated *ingenious*, *fascinating* and *jaw dropping* with greater probability. This identifies some classic instances of bias in data arising from social norms and structures. However, not all bias observed in data are against the presumed *unprivileged* community, for example, speakers from other genders get higher probability for *courageous* label as compared to male speakers. Our goal is to remove all types of bias from data, both expected and unexpected. The counterfactual fairness is agnostic to the type of bias and aims to remove all possible unfairness in rating across all possible combinations of sensitive attributes. Moreover it can be shown that under suitable assumptions, counterfactual fairness implies group fairness (see arxiv version for details).

6 FairyTED Pipeline

To achieve the goal of building a fair predictor of TED talk ratings Y , we execute the following steps:

1. **Show bias in real data:** We show that the viewer ratings in the TED talks is actually biased by using SPD and DI. This justifies the need to build a fair classification model for TED talk ratings.
2. **Preprocess data and define causal model:** We then define a causal model, C-DAG of the TED talk ratings, which consist of three major components: the unobserved variables U , the sensitive attributes S and the TED talk video attributes $X = \{T, V, Y\}$ obtained by preprocessing the data (see Section 4.2). For a specific model shown




Figure 1: Existence of bias in data: Here we show an example of biased rating in data using disparate SPD and DI metrics (see Observation in Data). In this example we compare male speakers with others and find male speakers are rated to give fascinating, ingenious and jaw dropping talks with higher probability. However contrary to expectation, other speakers had higher probability of being courageous.

in Fig. 3, U is the skill set and background of the speaker and S is gender and race. U and S causally influence T , V and Y . The skill set, background, gender and race of the speakers strongly influence their life experiences and hence govern the content of their talks. These also determine whether viewers choose to view their talk or not and what type of rating they get if viewed. Since other similar models can also be justified for the dataset, we ensure that our system is robust to any kind of causal model with similar setup (see Fig. 3).

- Model average:** We consider variants of C-DAG with two intuitively possible manipulations, 1) C-DAG1: unobserved causes affecting T and V are independent, meaning U decomposes into U_1 and U_2 such that U_1 affects T and U_2 affects V (e.g. skill set only influences how likely a talk will be viewed and background of the speaker influences the content of the talk, shown in Fig. 3(b)). 2) C-DAG2: the affect of sensitive attributes is manipulated, we consider the case where gender does not influence T (Fig. 3(c)).
- Fit model parameters:** We fit the parameters of each causal model, C-DAG, C-DAG1, C-DAG2.
- Create augmented data with counterfactual samples:** Next from each of these fitted models, we generate counterfactual samples of S (such as replacing male speaker with female speaker for a particular talk with fixed skill set and background, also see Preliminaries) and create augmented datasets (D_{aug} , $D1_{aug}$, $D2_{aug}$) to be used for classification.
- Train fair classifier:** Finally for each model we use corresponding augmented dataset to train a neural network for binary classification of each of the 14 rating labels. The loss function used to train the network has two parts: 1) the first part minimizes prediction error when compared to true data labels and 2) the second part reduces disparity between the labels of observed values of S and their

corresponding counterfactuals. This ensures that simply changing a male speaker to female with fixed skill set and background does not influence the rating. The prediction accuracy of the fair classifier is obtained by averaging performance across all three models.

- Fairness validation:** We finally validate that our classifier is counterfactually fair as compared to actual ratings provided by viewers. For this we introduce a novel metric *coefficient of probability variance*, CV_{prob} that compares variability of ratings across possible instances of S before and after introduction of fairness measure (e.g. hypothetically if male and female speakers were rated funny with probability 0.75 and 0.45 just due to difference in S and after fairness was introduced in prediction this variability dropped, becoming 0.75 and 0.70 solely based on same content, skill set and background).

This pipeline brings together fairness measuring metrics and counterfactual fairness incorporation techniques to build a complete setup for TED talk dataset (Fig. 2). Our setup can be applied to any language or video dataset whose feature embedding can be obtained using any state-of-the-art method. In addition, this setup can accommodate multiple causal models to ensure fairness in classifier across all possible models. Besides this, our setup also allows having nodes that cater to unobserved causes in the world.

6.1 Causal Model

We consider three relevant variants of causal models for TED talk data using the general definitions mentioned in preliminaries.

- C-DAG as in Fig. 3(a),

$$\begin{aligned}
 T &\sim \mathcal{N}(w_T^U U + w_T^S S, \sigma_T^2 \mathbf{1}) \\
 V &\sim \mathcal{N}^{[0,1]}(\sigma(w_V^U U + w_V^S S + w_V^T T), \sigma_V^2 \mathbf{1}) \\
 Y &\sim \text{Bern}(\sigma(w_Y^U U + w_Y^S S + w_Y^T T + w_Y^V V))
 \end{aligned}$$

where $\mathcal{N}^{[0,1]}$ denotes the truncated Gaussian in the domain $[0, 1]$. We generate U from $\mathcal{N}(0, \mathbf{1})$ and fit $\Theta_{\text{C-DAG}} = \{w_T^U, w_T^S, \sigma_T, w_V^U, w_V^S, w_V^T, \sigma_V, w_Y^U, w_Y^S, w_Y^T, w_Y^V, \sigma_Y\}$ using variational inference algorithm in PyMC3 (Salvatier, Wiecki, and Fonnesebeck 2016).

- We modify C-DAG such that U decomposes into mutually exclusive sets U_1 and U_2 to influence T and V independently (Fig. 3(b)) giving C-DAG1,

$$\begin{aligned}
 T &\sim \mathcal{N}(w_T^{U_1} U_1 + w_T^S S, \sigma_T^2 \mathbf{1}) \\
 V &\sim \mathcal{N}^{[0,1]}(\sigma(w_V^{U_2} U_2 + w_V^S S + w_V^T T), \sigma_V^2 \mathbf{1}) \\
 Y &\sim \text{Bern}(\sigma(w_Y^{U_1} U_1 + w_Y^S S + w_Y^T T + w_Y^V V))
 \end{aligned}$$

Both U_1 and U_2 are drawn from $\mathcal{N}(0, \mathbf{1})$ and we fit $\Theta_{\text{C-DAG1}} = \{w_T^{U_1}, w_T^S, \sigma_T, w_V^{U_2}, w_V^S, w_V^T, \sigma_V, w_Y^{U_1}, w_Y^S, w_Y^T, w_Y^V, \sigma_Y\}$ to obtain posterior distributions over U_1 and U_2 .




Figure 2: FairyTED pipeline: The setup to build a fair classifier for TED talk dataset (see Section 6). First step is to confirm presence of bias in data. Next we preprocess the data to obtain relevant attributes (details in Section 4). Following that we build possible causal models for the preprocessed data and generate augmented datasets with counterfactual samples (see Section 6.1). We then train a classifier for the augmented datasets with a loss function incorporating a fairness measure (see Section 6.2). As a final step we validate that our system learns to generate fair prediction of ratings for the TED talk dataset (see Results).

- In C-DAG2 S from C-DAG decomposes into S_1 and S_2 to influence T and V as in Fig. 3(c). $S_1 \subset S$ consists of race only, whereas, $S_2 = S$ which includes both race and gender,

$$T \sim \mathcal{N}(w_T^U U + w_T^{S_1} S_1, \sigma_T^2 \mathbf{1})$$

$$V \sim \mathcal{N}^{[0,1]}(\sigma(w_V^U U + w_V^{S_2} S_2 + w_V^T T), \sigma_V^2 \mathbf{1})$$

$$Y \sim \text{Bern}(\sigma(w_Y^U U + w_Y^S S + w_Y^T T + w_Y^V V))$$

Here, $U \sim \mathcal{N}(0, \mathbf{1})$ and we fit $\Theta_{\text{C-DAG2}} = \{w_T^U, w_T^{S_1}, \sigma_T, w_V^U, w_V^{S_2}, w_V^T, \sigma_V, w_Y^U, w_Y^S, w_Y^T, w_Y^V, \sigma_Y\}$ to obtain posterior distribution over U .

From each of the three models we generate counterfactual samples of S .

6.2 Classifier Model

We train a neural network with one hidden layer of 400 nodes to predict ratings Y . Let $\{(s_i, x_i)\}_{i=1}^N, \{y^i\}_{i=1}^N$ represent the attributes and labels of the dataset where s_i represents an instance of S and x_i represents instance of (T, V) . We train a classification function g such that $\hat{Y} = g(s, x)$ using a loss function which is a combination of *prediction loss* and *unfairness loss*. We use binary cross entropy loss (BCE) to calculate the prediction error and an unfairness function u to estimate the unfairness of the classifier as,

$$\mathcal{L}(g) = \frac{1}{N} \sum_{i=1}^N \left(\underbrace{\text{BCE}(g(s_i, x_i), y^i)}_{\text{prediction loss}} + \gamma \underbrace{\sum_{s' \neq s_i} u(g, s', s_i, x_i)}_{\text{unfairness loss}} \right) \quad (1)$$

$$u(g, s', s_i, x_i) = \frac{1}{C} \sum_{c=1}^C \max\{0, |g(s_i, x_i^c) - g(s'_i, x_i^c)| - \epsilon\} \quad (2)$$

where C represents the number of counterfactual samples for each observed data instance and ϵ is a hyperparameter which makes sure that our predictor maintains a (ϵ, δ) -approximate counterfactual fairness (δ is a function of γ , for more details about the choice of the unfairness function, please refer to (Russell et al. 2017)). We tune γ and ϵ to obtain best results in our causal models, see Table 3.

7 Results

7.1 Prediction Accuracy

After training the classifier with augmented datasets for the three models, we obtain an average prediction accuracy of 69% across all rating labels (Table 3). This accuracy is obtained by training the classifier without the unfairness measure in the loss function. The mean accuracy obtained from an unfair classifier is comparable to accuracy reported in recent studies on TED talk data (Cheng et al. 2014; Tanveer et al. 2019). However, note that our language model (Le and Mikolov 2014) is much simpler when compared to methods used in the cited studies. With this simple choice we emphasize the general appeal of our approach whose goal is to reduce unfair prediction in data irrespective of the embedding technique. After addition of the unfairness measure u in the loss function in equation (1), the average prediction accuracy goes down as expected, to 67% but not significantly (see section 6.2 for details). The hyperparameter γ in the loss function plays a critical role in determining the trade-off between fairness in prediction and its accuracy, the smaller the value of γ , the more unfair the prediction.




Figure 3: Causal models: Here we show three causal models considered for our FairyTED setup. C-DAG consists of all relevant dependencies between unobserved attributes U , sensitive attributes S , data attribute V (view) and rating prediction Y . In C-DAG1 we manipulate the influence of U on V and in C-DAG2 we manipulate the influence of S on V . This setup allows averaging across possible causal models when the true model is unknown.

7.2 Fairness Improvement

We first show that there is a decrease in the unfairness measure of the classifier with the increase in training iterations (Fig 4(A)). We then verify the improvement of fairness in prediction of all 14 rating labels across possible groups of S . To do so we come up with a fairness comparison metric CV_{prob} and compare prediction fairness before and after addition of unfairness function u in the loss function in equation (1).

Table 3: The accuracy of the classifier is reported for different set of hyperparameters (ϵ, γ) . Note that the accuracy increases if γ is reduced or ϵ is increased (increase unfairness). When $(\epsilon, \gamma) = (\infty, 0)$ (most unfair) it achieves the highest accuracy.

(ϵ, γ)	(0.1,0.2)	(0.5,0.2)	0.1,0.4)	$(\infty, 0)$
beautiful	0.7	0.7	0.66	0.77
confusing	0.63	0.62	0.6	0.66
courageous	0.71	0.72	0.69	0.74
fascinating	0.68	0.62	0.64	0.7
funny	0.71	0.67	0.67	0.71
informative	0.71	0.66	0.65	0.7
ingenious	0.7	0.65	0.65	0.7
inspiring	0.66	0.68	0.61	0.72
jaw-dropping	0.64	0.61	0.58	0.66
longwinded	0.64	0.57	0.6	0.61
obnoxious	0.58	0.59	0.59	0.62
ok	0.62	0.59	0.57	0.66
persuasive	0.71	0.67	0.66	0.75
unconvincing	0.62	0.62	0.63	0.66

7.3 CV_{prob} Metric

We have 3 types of gender (male, female and other) and 4 types of race (White, Asian, African American and other) under S giving 12 possible groups denoted by $G = \{G_1, G_2, \dots, G_{12}\}$ on whom counterfactual fairness is tested. We then calculate with what probability each of these 12 groups obtain a particular rating, i.e., for all $k \in \{1, \dots, 14\}$ and $i \in \{1, \dots, 12\}$, we compute $p_i^k =$

$\mathbb{P}(y_k = 1 | S = G_i)$. We denote, $P^k = (p_1^k, \dots, p_{12}^k)$. Similarly for the predicted label we compute $\hat{p}_i^k = \mathbb{P}(\hat{y}_k = 1 | S = G_i)$ and denote, $\hat{P}^k = (\hat{p}_1^k, \dots, \hat{p}_{12}^k)$. Note that, variability in the coordinates of P^k is a measure of fairness across the group for the rating y_k . In particular, the more variable these coordinates are, the more unfair the label is. Also $CV = \frac{std}{mean}$ is a common statistical metric used to quantify variability/irregularity in a set of values. Hence we define, CV_{prob}^k as CV for the coordinates of P^k . Similarly, $\hat{C}V_{prob}^k$ is the CV for the coordinates of \hat{P}^k . Hence, if a predictor improves fairness in the prediction, then $\hat{C}V_{prob}^k$ should be less than CV_{prob}^k as in Fig. 4(B). We also compared the SPD and DI for the true labels and predicted labels of the test dataset (Fig. 4 (C,D)).

8 Conclusion

The FairyTED setup is applicable to any dataset with the following properties, 1) contains sensitive attributes which can cause biased predictions, 2) it is possible to define a causal model with relevant attributes of the dataset, 3) a counterfactual fairness measure can be defined on the prediction probability by using counterfactual samples of the sensitive attributes.

We successfully identified the above properties in the TED talk dataset and removed bias/unfairness in rating prediction. The impact of this work is many-fold: 1) First, it identifies the necessity of applying counterfactual fairness on the rich and influential database of TED talk videos. 2) Second, it identifies that counterfactual fairness measure is the most relevant for TED talk videos and like datasets. This is because it allows identification of attributes in data which play a critical role in causing biased predictions for example male versus female speakers. In public speaking platforms it is expensive and implausible to test the change in rating of a talk when given by a female speaker instead of a male speaker with the same content and same skill set. Counterfactual fairness allows to hypothetically test any such example and correct for the resulting bias. 3) In most real world situations it is hard to know the true causal model for the observed dataset, and our setup can

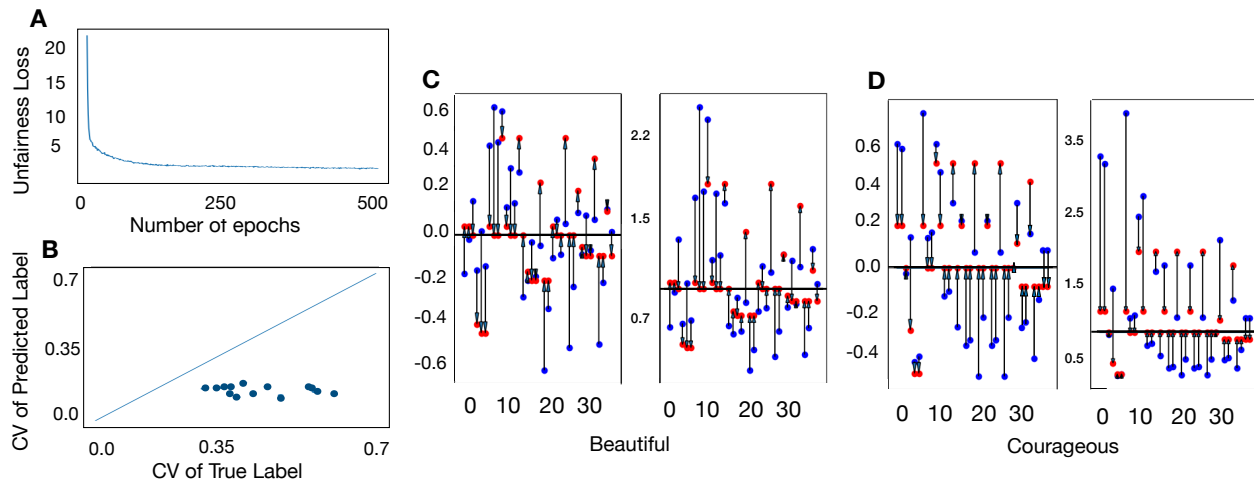


Figure 4: Increase in fairness of classifier: (A) The unfairness measure of the classifier decreases with the increase in training iterations (B) CV_{prob} values of true vs predicted labels. All points fall under the $y = x$ diagonal line, implying CV_{prob} is reduced by the fair classifier. (C) SPD moves towards 0 and (D) DI moves towards 1 for most of all possible pairs of groups across S after the prediction. The blue dots denote the fairness measures for the true labels and red dots denote the fairness measures for the predicted labels

deal with this issue as long as we have some idea of possible models. 4) In rich and complicated domain such as TED talks, there can be a lot of unobserved attributes that can affect the data and our setup can take care of it by inferring a posterior distribution over the unobserved attributes. 5) Any dataset with a simple embedding scheme can use this model. One straightforward extension of this framework can be on job interview based datasets. We can include better encoding schemes such as (Devlin et al. 2018; Yang et al. 2019) to obtain even better prediction accuracy besides making it fair. We can also include multi-modal (e.g. audio-visual) information from the dataset to obtain rich representation and test how counterfactual fairness measure generalizes across various modes. 6) We propose an intuitive novel metric that quantifies degree of fairness employed by our setup. Besides these, information from temporal evolution can also be used to improve our framework in future work. Finally the FairyTED setup has three important social impacts: 1) It ensures that speakers get fair feedback and can improve only based on fair fallacies 2) Organizers can employ diverse speakers without worrying about degradation in rating when skill, ability, influence and content are matched. For example, suppose a talk on global warming was given by a male speaker and obtained good rating. With our setup, the organisers can choose a female speaker with comparable qualities as the male to speak on the global warming without worrying about deterioration in rating. 3) Finally, new viewers will not be biased by prevalent unwanted biased ratings from previous users and prevents propagation of unfairness over time.

9 Acknowledgments

We thank the anonymous reviewers for their valuable suggestions.

References

- Alikaniotis, D.; Yannakoudakis, H.; and Rei, M. 2016. Automatic text scoring using neural networks. *arXiv preprint arXiv:1606.04289*.
- Baeza-Yates, R. 2018. Bias on the web. *Communications of the ACM* 61(6):54–61.
- Bellamy, R. K. E.; Dey, K.; Hind, M.; Hoffman, S. C.; Houde, S.; Kannan, K.; Lohia, P.; Martino, J.; Mehta, S.; Mojsilovic, A.; Nagar, S.; Ramamurthy, K. N.; Richards, J.; Saha, D.; Sattigeri, P.; Singh, M.; Varshney, K. R.; and Zhang, Y. 2018. AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias.
- Bickel, P. J.; Hammel, E. A.; and O’Connell, J. W. 1975. Sex bias in graduate admissions: Data from berkeley. *Science* 187(4175):398–404.
- Biddle, D. 2006. *Adverse impact and test validation: A practitioner’s guide to valid and defensible employment testing*. Gower Publishing, Ltd.
- Bolukbasi, T.; Chang, K.-W.; Zou, J. Y.; Saligrama, V.; and Kalai, A. T. 2016. Man is to computer programmer as woman is to homemaker? debiasing word embeddings. In *Advances in neural information processing systems*, 4349–4357.
- Brennan, T.; Dieterich, W.; and Ehret, B. 2009. Evaluating the predictive validity of the compas risk and needs assessment system. *Criminal Justice and Behavior* 36(1):21–40.
- Buolamwini, J., and Gebru, T. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*, 77–91.
- Calders, T., and Verwer, S. 2010. Three naive bayes ap-

- proaches for discrimination-free classification. *Data Mining and Knowledge Discovery* 21(2):277–292.
- Chen, L.; Zhao, R.; Leong, C. W.; Lehman, B.; Feng, G.; and Hoque, M. E. 2017a. Automated video interview judgment on a large-sized corpus collected online. In *2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII)*, 504–509. IEEE.
- Chen, L.; Zhao, R.; Leong, C. W.; Lehman, B.; Feng, G.; and Hoque, M. E. 2017b. Automated video interview judgment on a large-sized corpus collected online. In *ACII*, 504–509.
- Cheng, D. S.; Salamin, H.; Salvagnini, P.; Cristani, M.; Vinciarelli, A.; and Murino, V. 2014. Predicting online lecture ratings based on gesturing and vocal behavior. *Journal on Multimodal User Interfaces* 8(2):151–160.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226. ACM.
- Garg, S.; Perot, V.; Limtiaco, N.; Taly, A.; Chi, E. H.; and Beutel, A. 2019. Counterfactual fairness in text classification through robustness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 219–226. ACM.
- Grgic-Hlaca, N.; Zafar, M. B.; Gummadi, K. P.; and Weller, A. 2016. The case for process fairness in learning: Feature selection for fair decision making. In *NIPS Symposium on Machine Learning and the Law*, volume 1, 2.
- Joseph, M.; Kearns, M.; Morgenstern, J.; Neel, S.; and Roth, A. 2016. Rawlsian fairness for machine learning. *arXiv preprint arXiv:1610.09559* 1(2).
- Kamiran, F., and Calders, T. 2009. Classifying without discriminating. In *2009 2nd International Conference on Computer, Control and Communication*, 1–6. IEEE.
- Kamishima, T.; Akaho, S.; and Sakuma, J. 2011. Fairness-aware learning through regularization approach. In *2011 IEEE 11th International Conference on Data Mining Workshops*, 643–650. IEEE.
- Kay, M.; Matuszek, C.; and Munson, S. A. 2015. Unequal representation and gender stereotypes in image search results for occupations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 3819–3828. ACM.
- Krippendorff, K. 2011. Agreement and information in the reliability of coding. *Communication Methods and Measures* 5(2):93–112.
- Kusner, M. J.; Loftus, J.; Russell, C.; and Silva, R. 2017. Counterfactual fairness. In *Advances in Neural Information Processing Systems*, 4066–4076.
- Le, Q., and Mikolov, T. 2014. Distributed representations of sentences and documents. In *International conference on machine learning*, 1188–1196.
- Mehrabi, N.; Morstatter, F.; Saxena, N.; Lerman, K.; and Galstyan, A. 2019. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*.
- Naim, I.; Tanveer, M. I.; Gildea, D.; and Hoque, M. E. 2016. Automated analysis and prediction of job interview performance. *IEEE Transactions on Affective Computing* 9(2):191–204.
- Nguyen, L. S., and Gatica-Perez, D. 2016. Hirability in the wild: Analysis of online conversational video resumes. *IEEE Transactions on Multimedia* 18(7):1422–1437.
- Pearl, J., et al. 2009. Causal inference in statistics: An overview. *Statistics surveys* 3:96–146.
- Russell, C.; Kusner, M. J.; Loftus, J.; and Silva, R. 2017. When worlds collide: integrating different counterfactual assumptions in fairness. In *Advances in Neural Information Processing Systems*, 6414–6423.
- Salvatier, J.; Wiecki, T. V.; and Fonnesbeck, C. 2016. Probabilistic programming in python using PyMC3. *PeerJ Computer Science* 2:e55.
- Schumann, C.; Wang, X.; Beutel, A.; Chen, J.; Qian, H.; and Chi, E. H. 2019. Transfer of machine learning fairness across domains. *arXiv preprint arXiv:1906.09688*.
- Taghipour, K., and Ng, H. T. 2016. A neural approach to automated essay scoring. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1882–1891.
- Tanveer, M. I.; Hasan, M. K.; Gildea, D.; and Hoque, M. E. 2019. A causality-guided prediction of the ted talk ratings from the speech-transcripts using neural networks. *arXiv preprint arXiv:1905.08392*.
- Tanveer, M. I.; Liu, J.; and Hoque, M. E. 2015. Unsupervised extraction of human-interpretable nonverbal behavioral cues in a public speaking scenario. In *23rd Annual ACM Conference on Multimedia*, 863–866. ACM.
- Tollenaar, N., and Van der Heijden, P. 2013. Which method predicts recidivism best?: a comparison of statistical, machine learning and data mining predictive models. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 176(2):565–584.
- Wilson, B.; Hoffman, J.; and Morgenstern, J. 2019. Predictive inequity in object detection. *arXiv preprint arXiv:1902.11097*.
- Yang, Z.; Dai, Z.; Yang, Y.; Carbonell, J.; Salakhutdinov, R.; and Le, Q. V. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*.
- Zafar, M. B.; Valera, I.; Gomez Rodriguez, M.; and Gummadi, K. P. 2017. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *Proceedings of the 26th International Conference on World Wide Web*, 1171–1180. International World Wide Web Conferences Steering Committee.
- Zliobaite, I. 2015. A survey on measuring indirect discrimination in machine learning. *arXiv preprint arXiv:1511.00148*.