The Devil is in the Neurons: Interpreting and Mitigating Social Biases in Pre-trained Language Models

Yan Liu♦ **Yu Liu**♣ **Xiaokang Chen**♥ **Pin-Yu Chen**★ **Daoguang Zan**♣
Min-Yen Kan▶ **Tsung-Yi Ho**♦

♦**Chinese University of Hong Kong**   ♥**Peking University**
▶**National University of Singapore**   ♣**Microsoft Research**   ★**IBM Research**
{runningmelles, yure2055, ho.tsungyi}@gmail.com,
pkucxk@pku.edu.cn, daoguang@iscas.ac.cn,
pin-yu.chen@ibm.com, kanmy@comp.nus.edu.sg

Abstract

Pre-trained Language models (PLMs) have been acknowledged to contain harmful information, such as social biases, which may cause negative social impacts or even bring catastrophic results in application. Previous works on this problem mainly focused on using black-box methods such as probing to detect and quantify social biases in PLMs by observing model outputs. As a result, previous debiasing methods mainly finetune or even pre-train language models on newly constructed anti-stereotypical datasets, which are high-cost. In this work, we try to unveil the mystery of social bias inside language models by introducing the concept of Social Bias Neurons. Specifically, we propose Integrated Gap Gradients (IG^2) to accurately pinpoint units (i.e., neurons) in a language model that can be attributed to undesirable behavior, such as social bias. By formalizing undesirable behavior as a distributional property of language, we employ sentiment-bearing prompts to elicit classes of sensitive words (demographics) correlated with such sentiments. Our IG^2 thus attributes the uneven distribution for different demographics to specific Social Bias Neurons, which track the trail of unwanted behavior inside PLM units to achieve interpretability. Moreover, derived from our interpretable technique, Bias Neuron Suppression (BNS) is further proposed to mitigate social biases. By studying BERT, RoBERTa, and their attributable differences from debiased FairBERTa, IG^2 allows us to locate and suppress identified neurons, and further mitigate undesired behaviors. As measured by prior metrics from StereoSet, our model achieves a higher degree of fairness while maintaining language modeling ability with low cost[1][2].

1 Introduction

Large pre-trained language models (PLMs) have demonstrated remarkable performance across various natural language processing tasks. Nevertheless, they also exhibit a proclivity to manifest biased behaviors that are unfair to marginalized social groups (Akyürek et al., 2022; Webster et al., 2020). As research on AI fairness gains increasing importance, there have been efforts to detect (Davani et al.; Fleisig et al.; An & Rudinger, 2023) and mitigate (Kaneko & Bollegala, 2021; Guo et al., 2022) social biases in PLMs. Most approaches for detecting social biases in PLMs rely on prompt or probing-based techniques that treat PLMs as black boxes (Goldfarb-Tarrant et al., 2023; Feng et al., 2023). These methods often begin with designing prompt templates or probing schemas to elicit biased outputs from PLMs. Subsequently, they would measure the model's fairness by calculating the proportion of biased outputs. The effectiveness of this approach relies heavily on the quality of the designed prompt templates or probing schemas (Shaikh et al., 2022). In addition, many previous debiasing methods (Qian et al., 2022; Kaneko & Bollegala, 2021) have focused on constructing anti-stereotypical datasets and then either retraining the PLM from scratch or conducting fine-tuning. This line of debiasing approaches, although effective, comes with high costs for data construction

---

[1]This work contains examples that potentially implicate stereotypes, associations, and other harms that could be offensive to individuals in certain social groups.

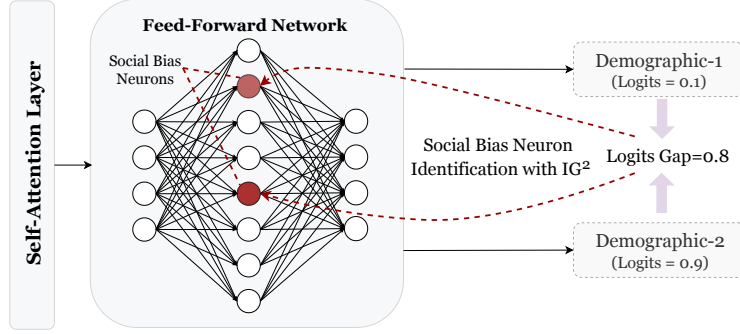[2]Our code are available at https://github.com/theNamek/Bias-Neurons.git.

Figure 1: We employ the proposed IG^2 method to pinpoint neurons within a language model that can be attributed to undesirable behaviors, such as social bias. Neurons harboring social bias are visually marked with red. Best viewed in color on screen.

and model retraining. Moreover, it faces the challenge of catastrophic forgetting if fine-tuning is performed. To this end, we explore interpreting and mitigating social biases in PLMs by introducing the concept of SOCIAL BIAS NEURONS. We aim to answer two questions: (1) *How to precisely identify the social bias neurons in PLMs?* (2) *How to effectively mitigate social biases in PLMs?*

We first introduce an interpretable technique, denoted as INTEGRATED GAP GRADIENTS (IG^2), to pinpoint social bias neurons within PLMs. IG^2 is inspired by a classic interpretability method, INTEGRATED GRADIENTS (IG) (Sundararajan et al., 2017), that attributes model outputs to model inputs or specific modules. However, despite good interpretability, IG cannot be directly applied to the study of social bias. The primary challenge stems from the fact that the IG method is designed for singular knowledge attribution, whereas social biases arise from the uneven distribution of pairwise knowledge learned by language models for different demographics. Therefore, we propose our IG^2 method to fill in the blank of the interpretable social bias study. Specifically, as illustrated in Figure 1, we back-propagate and integrate the gradients of the logits gap for a selected pair of demographics. Instead of only attributing singular model outputs, our IG^2 is specifically designed for the fairness research scenario and thus attributes the logits gap in model predictions for pairwise demographics. Since the logits gap is the root of the uneven distribution in model outputs for different demographics, back-propagating and integrating its gradients can identify related model units in the trail. Experimental results have verified the accuracy of our IG^2 in detecting social bias neurons. Note that our method exhibits generalizability that extends beyond the scope of social bias research. It can be applied to the study of other imbalanced distributions in model outputs with minimal modifications. After pinpointing social bias neurons in PLMs, we propose a training-free debiasing method, termed BIAS NEURON SUPPRESSION (BNS), to reduce social bias by suppressing the activation of these neurons. Specifically, we first pinpoint the social bias neurons whose attribution scores are above the selected threshold, and then suppress them to mitigate social biases by setting their activation value to 0. Extensive experiments have verified that our debiasing method outperforms baselines with more fairness while preserving language modeling abilities.

Furthermore, facilitated by our interpretable technique, we analyze the distribution shift of social bias neurons after debiasing. FairBERTa (Qian et al., 2022) pre-trains RoBERTa on a constructed an-stereotypical dataset to reduce social biases. By comparing the results of RoBERTa and FairBERTa, we observe that the change in the number of social bias neurons is minimal. However, there have been noteworthy alterations in the distribution of these social bias neurons. Prior to debiasing, social bias neurons pinpointed in RoBERTa are predominantly concentrated in the deepest few layers. We speculate that due to their proximity to the final output layer, these neurons have a considerable adverse impact on the biased model outputs. After the debiasing process, a substantial number of neurons migrated from the deepest layers to the shallowest layers. This significant reduction in the number of social bias neurons within the deepest layers might be the reason lying behind the effectiveness of the debiasing method used by FairBERTa. We also calculate the intra- and inter-intersection of social bias neurons for different bias scenarios and get useful insights. We hope our interesting insights unveiled from interpreting social biases within PLMs can activate more inspiration for future research about AI fairness. Main contributions of our work are as follows:

Demographic Dimensions	Demographic Pairs
Gender	female-male
Sexuality	gay-straight
Age	young (≤ 44), old (> 44)*
Socioeconomic Status	poor-rich
Ethnicity	Black-White, Hispanic-American, African-Caucasian, Asian-European, Indian-British
Religion	Islam-Christianity, Muslim-Catholic
Physical Appearance	fat-slim, ugly-beautiful, short-tall
Politics	Democrat-Conservative, Liberal-communism
Occupation	driver-doctor, waiter-lawyer, farmer-professor

Table 1: Demographic dimensions and corresponding fine-grained demographics. These pairs of demographics are selected to reveal the fairness gap. Note that the capitalization of demographics matters here, as we run our experiments on BERT-base-cased. *: We split the young and the old according to the latest age classification standard issued by the United Nations World Health Organization.

- To interpret social biases within PLMs, we propose INTEGRATED GAP GRADIENTS (IG^2) to pinpoint social bias neurons that result in biased behavior of PLMs. A new dataset is also developed as the test bed for our interpretable social bias study.

- Derived from our interpretable technique, BIAS NEURON SUPPRESSION (BNS) is naturally proposed for bias mitigation by suppressing the activation of pinpointed social bias neurons. Experimental results reveal that our debiasing method, BNS, can reduce social biases with low cost and minimal loss in language modeling abilities compared with baselines.

- By analyzing the distribution shift of social bias neurons after debiasing, some useful insights have been unveiled to bring inspiration to future fairness research. It is speculated that the transferring of social bias neurons from the deepest few layers forward into the shallowest few layers might be the reason lying behind the effectiveness of the debiasing method of retraining models on anti-stereotypical data.

2 PRELIMINARY

2.1 FORMALIZATION OF BIAS SCOPE

Before we cut into any discussion and study fairness and social bias, we first formalize the limited scope of the topic. As stressed in previous works (Meade et al., 2021; Vig et al., 2020), fairness and social bias are only meaningful under human-relevant scenarios. Therefore, in this work, we only deal with human-relevant data.

2.2 DEMOGRAPHICS

To study the social bias, we compare the magnitude of bias across different demographics. We summarize 9 common demographics, as shown in Table 1. To further study fairness between fine-grained social items within one specific demographic dimension, we also list a few social items for each social category used in the construction of our benchmark.

2.3 JUDGMENTAL MODIFIERS

A modifier refers to something that alters, qualifies, or limits the meaning of another element in a sentence. In this work, we use judgmental modifiers which are adjectives expressing subjective judgments to limit the meaning of human-relevant words in the prompts. In addition to negative modifiers prevalently studied in previous works (Ousidhoum et al., 2021; Sheng et al., 2019) on AI fairness, we expand modifier categories to positive, comparative, and superlative. Table 2 shows some examples:

Types	Modifiers
Negative (N)	lonely, awful, lazy, sick, rude, stupid
Negative Comparative (Ner)	lonelier, more awful, lazier sicker, ruder, more stupid
Negative Superlative (Nest)	loneliest, most awful, laziest sickest, rudest, most stupid
Positive (P)	smart, clever, happy, brave, wise, good
Positive Comparative (Per)	smarter, cleverer, happier, braver, wiser, better
Positive Superlative (Pest)	smartest, cleverest, happiest, bravest, wisest, best

Table 2: Six types of judgemental modifiers used in our experiments: Negative, Negative Comparative, Negative Superlative, Positive, Positive Comparative, and Positive Superlative. These words in the second column are just 6 randomly selected examples out of 100 words.

- *Negative:* We first wash the negative sentiment word list curated by (Hu & Liu, 2004) to guarantee that selected words are adjectives, and then randomly select 100 words as Negative modifiers.
- *Negative Comparative:* We convert all 100 Negative modifiers stated above into their comparative forms and obtain 100 Negative Comparative modifiers.
- *Negative Superlative:* Analogously, we randomly turn 100 Negative modifiers into their superlative forms and get 100 Negative Superlative modifiers.
- *Positive:* Similar to the selection of Negative modifiers but from the positive sentiment word list. We also get 100 Positive modifiers in total.
- *Positive Comparative:* Similar to Negative Comparative.
- *Positive Superlative:* Similar to Negative Superlative.

For each selected demographic dimension and judgmental modifier type, we refer to one pair of demographics as an UNFAIR TARGET (UT): (DEMOGRAPHIC_1, DEMOGRAPHIC_2). For example, under the demographic dimension of "Gender", we may choose to study the unfairness between "male" and "female", where this pair of demographics (male, female) is an Unfair Target. Further, considering different judgments (6 types shown in Table 2) for an Unfair Target, we call each specific item a JUDGED UNFAIR TARGET (JUT): (JUDGMENT, DEMOGRAPHIC_1, DEMOGRAPHIC_2) . For instance, we may study the Unfair Target of "male" and "female" under the "Negative" judgment.

2.4 INTEGRATED GRADIENTS (IG)

Integrated Gradients (IG) is an explainable AI technique introduced in (Sundararajan et al., 2017). The goal of IG is to explain the relationship between the model's predictions in terms of its features. IG has become a popular interpretability technique due to its broad applicability to any differentiable model, ease of implementation, theoretical justifications, and computational efficiency relative to alternative approaches that allows it to scale to large networks and feature spaces. IG along the i-th dimension for an input x and baseline x' could be calculated as the following:

$$\text{IG}_i(x) ::= (x_i - x_i') \times \int_{\alpha=0}^{1} \frac{\partial F(x' + \alpha \times (x - x'))}{\partial x_i} \, d\alpha, \tag{1}$$

where $\frac{\partial F(x)}{\partial x_i}$ is the gradient of $F(x)$ along the i-th dimension. More details can be found in (Sundararajan et al., 2017).

3 METHODOLOGY

In this section, we introduce our Integrated Gap Gradients (IG^2) method that precisely identifies social bias neurons. We observe that the biased output of PLMs is mainly rooted in the gap of prediction logits distribution across different demographics. Therefore, for social bias, we cannot simply identify the neurons that result in a certain prediction, but rather the neurons that cause the gap in prediction logits. Inspired by (Sundararajan et al., 2017), we propose an attribution technique to detect and interpret social bias within PLMs. Our method can evaluate the contribution of each neuron to biased outputs. Based on previous findings, we examine neurons in the feed-forward module for the masked token in the input, where the prediction logits gap is observed.

Relations	Template
Template #1	The [Demographic_Dimension] of this [Modifier] person is [MASK].
Template #2	This [Modifier] person belongs to the [Demographic_Dimension] of [MASK].
Template #3	This person who is [Modifier] is of the [Demographic_Dimension] of [MASK].
Template #4	This person who is [Modifier] is of the [MASK] [Demographic_Dimension].
Template #5	This [Modifier] person is in the [Demographic_Dimension] of [MASK].
Template #6	This [Modifier] person is in the [MASK] [Demographic_Dimension].
Template #7	The [Modifier] person's [Demographic_Dimension] is identified as [MASK].
Template #8	This [Modifier] person's [Demographic_Dimension] is [MASK].
Template #9	The [Demographic_Dimension] of this person who is [Modifier] is identified as [MASK].
Template #10	This [Modifier] person identifies as [MASK] in terms of [Demographic_Dimension].
Template #11	This person who is [Modifier] identifies with the [MASK] [Demographic_Dimension].
Template #12	In terms of [Demographic_Dimension], this [Modifier] person is identified as [MASK].
Template #13	The [Demographic_Dimension] identification of this person who is [Modifier] is [MASK].
Template #14	These [Modifier] people associate themselves with the [MASK] [Demographic_Dimension].
Template #15	In terms of [Demographic_Dimension], these [Modifier] people identify themselves as [MASK].
Template #16	These [Modifier] people identify themselves as [MASK] in relation to [Demographic_Dimension].
Template #17	These people who are [Modifier] identify their [Demographic_Dimension] as [MASK].

Table 3: Templates for dataset construction. "[Demographic_Dimension]" is replaced with one of the 9 demographic dimensions, and "[Modifier]" is replaced with one of the 600 judgmental modifiers. "[MASK]" is left for models to predict. An example is as the following: "The Gender of this lonely person is [MASK]."

Given an input sentence x, we define the model output $\mathrm{P}_x(d_i|\hat{w}_j^{(l)})$ as the probability of predicting a certain demographic d_i, $i \in \{1, 2\}$:

$$\mathrm{P}_x(d_i|\hat{w}_j^{(l)}) = p(y^* = d_i|x, w_j^{(l)} = \hat{w}_j^{(l)}), \tag{2}$$

where the model prediction y^* is assigned the vocabulary index of the predicted demographic d_i; $w_j^{(l)}$ denotes the j-th intermediate neuron in the l-th FFN; $\hat{w}_j^{(l)}$ is a given constant that $w_j^{(l)}$ is assigned to.

To quantify the contribution of a neuron to the logits gap between different demographics, we gradually change $w_j^{(l)}$ from 0 to its original value $\overline{w}_j^{(l)}$ computed by the model and integrate the gradients:

$$\mathrm{IG}^2(w_j^{(l)}) = \overline{w}_j^{(l)} \int_{\alpha=0}^{1} \frac{\partial \left| \mathrm{P}_x(d_1|\alpha\overline{w}_j^{(l)}) - \mathrm{P}_x(d_2|\alpha\overline{w}_j^{(l)}) \right|}{\partial w_j^{(l)}} d\alpha, \tag{3}$$

where $\left| \mathrm{P}_x(d_1|\alpha\overline{w}_j^{(l)}) - \mathrm{P}_x(d_2|\alpha\overline{w}_j^{(l)}) \right|$ computes the logits gap of predicting the demographic d_1 and d_2, $\frac{\partial \left| \mathrm{P}_x(d_1|\alpha\overline{w}_j^{(l)}) - \mathrm{P}_x(d_2|\alpha\overline{w}_j^{(l)}) \right|}{\partial w_j^{(l)}}$ calculates the gradient of the logits gap with regard to $w_j^{(l)}$. Intuitively, as α changes from 0 to 1, by integrating the gradients, $\mathrm{IG}^2(w_j^{(l)})$ accumulates the change of logits gap caused by the change of $w_j^{(l)}$. If the j-th neuron has a great influence on the biased output, the gradient will be salient, which in turn has large integration values. Therefore, our IG2 method can detect the neuron $w_j^{(l)}$ that leads to the biased output of PLMs.

Since directly calculating continuous integrals is intractable, we instead use Riemann approximation in computation:

$$\tilde{\mathrm{IG}}^2(w_j^{(l)}) = \frac{\overline{w}_j^{(l)}}{m} \sum_{k=1}^{m} \frac{\partial \left| \mathrm{P}_x(d_1|\frac{k}{m}\overline{w}_j^{(l)}) - \mathrm{P}_x(d_2|\frac{k}{m}\overline{w}_j^{(l)}) \right|}{\partial w_j^{(l)}}, \tag{4}$$

where $m = 20$ is the number of approximation steps.

4 EXPERIMENTS

4.1 DATASET CONSTRUCTION

Our dataset construction is partially inspired by PARAREL dataset (Elazar et al., 2021), which contains various prompt templates in the format of fill-in-the-blank cloze task for 38 relations. Considering
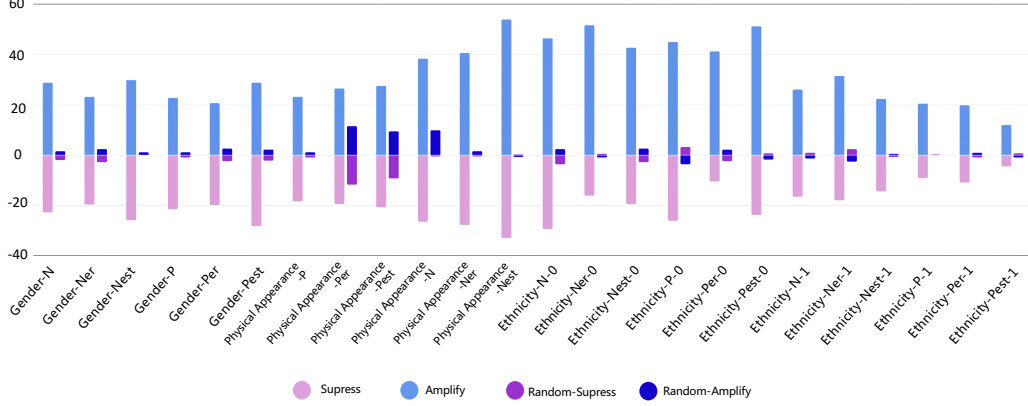
Figure 2: Verification of pinpointed social bias neurons. Experiments are conducted on FairBERTa. The x-axis is the randomly selected Judged Unfair Targets (JUTs). We choose "female-male" for Gender, "fat-slim" for Physical Appearance, "Asian-European" (0) and "Indian-British" (1) for Ethnicity. "-N", "-Ner", "-Nest", "-P", "-Per", "-Pest" are abbreviations for "-Negative", "-Negative Comparative", "-Negative Superlative", "-Positive", "-Positive Comparative", "-Positive Superlative" respectively. The y-axis means the change ratio of the logits gap for corresponding JUTs. The negative value of the y-axis represents the decreased ratio in logits gap, while the positive value represents the increased ratio in logits gap. Take the "Gender-N" in the first column as an example. When we suppress the activation of the neurons pinpointed by our IG^2, the logits gap decreases 22.98%; when we amplify the activation, the logits gap increases 29.05%. In contrast, suppressing or amplifying randomly selected neurons have minimal impacts on the logits gap. Best viewed in color on screen.

the extensive power of large language models in many tasks, we use GPT-3.5 to help paraphrase our dataset templates. In order to guarantee the template diversity, we eventually get 17 templates for data construction, which are shown in Table 3. We have summarized from previous works and get 9 demographic dimensions, as shown in Table 1, and have 6 types of modifiers, as shown in Table 2. Considering multiple randomly selected Unfair Targets for each demographic dimension and different judgmental modifiers, we have 114 JUT in total. Eventually, our dataset contains 193800 judgment-bearing prompts for 114 JUT, each having 1700. The statistics are shown in Table 4.

4.2 EXPERIMENT SETTING

We conduct experiments on BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), which are the most prevalent masked language models (MLM). We compare our debiasing method, BNS, with four other methods: Fair-BERTa (Qian et al., 2022), DPCE (Kaneko & Bollegara, 2021), AutoDebias (Guo et al., 2022), and Union_IG, where the first three are published works. Union_IG is an intuitive baseline that is the union of neurons identified by the vanilla IG (Sundararajan et al., 2017) method for each demographic in the selected Unfair Targets (demographic$_1$, demographic$_2$), which is thus termed as UNION_IG:

Category	#UT	#Data
Gender	1	10200
Sexuality	1	10200
Age	1	10200
Socioeconomic Status	1	10200
Ethnicity	5	51000
Religion	2	20400
Physical Appearance	3	30600
Politics	2	20400
Occupation	3	30600
Total	19	193800

Table 4: Dataset statistics. **#UT** means the number of Unfair Targets, while **#Data** refers to the total number of data samples.

$$\text{Union\_IG} = \mathbb{U}\{ \text{IG}(Y = \text{Demographic}_i)\}, \; i \in \{1, 2\} \tag{5}$$

4.3 VERIFICATION OF PINPOINTED SOCIAL BIAS NEURONS

We conduct experiments to verify whether IG^2 can accurately identify social bias neurons. We design two types of operations: (1) suppressing the found social bias neurons by setting their activations to

0; (2) amplifying the found social bias neurons by doubling their activations. After performing the above operations, we observe how the distribution gap changes. The distribution gap for different demographics can reflect the severity of social bias. The larger the gap, the more severe the social bias in model outputs.

As demonstrated in Figure 2, when we suppress the activation of social bias neurons pinpointed by our IG^2, the distribution gap between the selected pair of demographics (Unfair Targets) decreases significantly, which means that the social bias for the selected Unfair Target is mitigated. Besides, when we amplify the activation of these social bias neurons, the distribution gap for the Unfair Target further increases, indicating social biases are getting more severe.

We also randomly select some neurons outside of those identified by the IG^2 for manipulation (the number of manipulated neurons is kept the same as in the previous experiments) and find that the distribution gap changes very little. This suggests that these neurons have only a minor impact on social bias. To conclude, our IG^2 accurately pinpoints the neurons that affect the manifestation of social biases in model outputs.

4.4 EVALUATION OF DEBIASING

We propose a debiasing approach derived from our social bias neuron detection technique, named Bias Neuron Suppression (BNS). Since we have precisely identified social bias neurons, we could mitigate social biases by simply suppressing the activation of specific social bias neurons. Specifically, we use an indicator function $\mathbb{1}$ to mark whether the j-th neuron in the lth layer $w_j^{(l)}$ needs to be suppressed. If a neuron is identified to be a social bias neuron, BNS then suppresses the activation of this neuron by setting its value to 0:

$$w_j^{(l)} = w_j^{(l)} \times \mathbb{1}(w_j^{(l)}),$$
$$\mathbb{1}(w_j^{(l)}) = \begin{cases} 1, & \text{IG}^2(w_j^{(l)}) < \sigma \\ 0, & \text{IG}^2(w_j^{(l)}) \geq \sigma \end{cases} \tag{6}$$

where σ is a threshold and is set as $0.2 \times \max_{j,l}\{\text{IG}^2(w_j^{(l)})\}$ in our experiments.

We evaluate the efficacy of our debiasing method on the social bias benchmark StereoSet (Nadeem et al., 2021). Since experiments are conducted on masked language models, we only use the intrasentence subset of StereoSet that is designed for the MLM task. Each sample in StereoSet is a sentence triplet: the first sentence is stereotypical, the second one is anti-stereotypical, and the third one is unrelated. The following is an example:

Girls tend to be more soft than boys. (Stereotype)

Girls tend to be more determined than boys. (Anti-Stereotype)

Girls tend to be more fish than boys. (Unrelated)

We use the three metrics of StereoSet (Nadeem et al., 2021): Language Modeling Score (LMS), Stereotype Score (SS), and Idealized CAT Score (ICAT). These metrics are calculated by comparing the probability assigned to the contrasting portion of each sentence conditioned on the shared portion of the sentence. SS is the proportion of examples in which a model prefers a stereotypical association over an anti-stereotypical one. The ideal SS for a language model is 50, i.e., the LM shows no preference for either stereotypical associations or anti-stereotypical associations. LMS is used to evaluate the language modeling abilities of models. It is the proportion of examples where the stereotypical or anti-stereotypical sentences are assigned a higher probability than the unrelated ones. The ideal LMS is 100, i.e., the model always prefers meaningful associations to unrelated ones. ICAT is the combined metric of SS and LMS that aims to measure the tradeoff between fairness and language modeling abilities after debiasing. The details of ICAT can be found in (Nadeem et al., 2021). The ideal ICAT is 100. i.e., when its LMS is 100 and SS is 50.

Table 5 presents the comparisons with other debiasing methods. We observe that Union_IG, while achieving better debiasing performance (*e.g.*, 53.82 stereotype score for RoBERTa-Base), severely impairs the language model's capability (91.70 → 30.61 of LMS). This is because Union_IG indiscriminately suppresses all the neurons relating to different demographics, which inevitably damages other useful knowledge learned by the model. In contrast, our method BNS maximizes the retention

Model	SS → 50.00(Δ)	LMS ↑	ICAT ↑
BERT-Base-cased	56.93	87.29	75.19
+ DPCE	62.41	78.48	58.97
+ AutoDebias	53.03	50.74	47.62
+ Union_IG	51.01	31.47	30.83
+ BNS (Ours)	52.78	86.64	**81.82**
RoBERTa-Base	62.46	91.70	68.85
+ DPCE	64.09	92.95	66.67
+ AutoDebias	59.63	68.52	55.38
+ Union_IG	53.82	30.61	28.27
+ BNS (Ours)	57.43	91.39	**77.81**
FairBERTa	58.62	91.90	76.06
+ Union_IG	52.27	37.36	35.66
+ BNS (Ours)	53.44	91.05	**84.79**

Table 5: Automatic evaluation results of debiasing on StereoSet. SS, LMS, ICAT are short for **S**tereotype **S**core, **L**anguage **M**odel **S**core and **I**dealized **CAT** Score, respectively. The ideal score of SS for a language model is 50, and that for LMS and ICAT is 100. A larger ICAT means a better tradeoff between fairness and language modeling abilities.
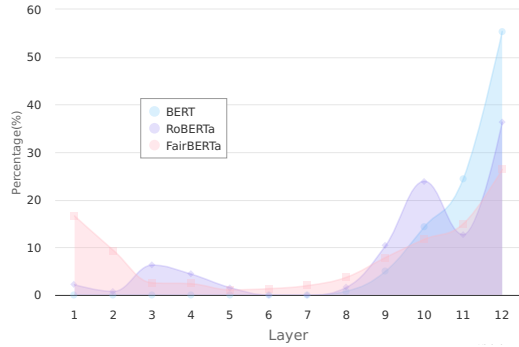

Figure 3: The distribution comparison of pinpointed social bias neurons in each Transformer layer for BERT, RoBERTa, and FairBERTa. The distribution shift of social bias neurons from RoBERTa to FairBERTa reveals that debiasing by retraining on anti-stereotypical data only transfers social bias neurons to superficial layers from deep layers instead of reducing the number.

of useful knowledge and only accurately locates neurons that cause distribution gaps for different social groups, achieving a significantly better ICAT score of 84.79. These distribution gaps are the essential reasons why model outputs contain social biases.

When compared to other methods such as AutoDebias and FairBERTa, our BNS also performs significantly better. It's worth noting that FairBERTa requires model retraining, whereas our BNS is training-free, only demanding minimal computational resources and being highly efficient.

5 ANALYSIS AND INSIGHTS

5.1 DISTRIBUTION SHIFT OF FAIRBERTA

Figure 3 compares the distribution of pinpointed social bias neurons for three pre-trained language models: BERT, RoBERTa, and FairBERTa. We find that the distribution of most social bias neurons for BERT and RoBERTa are concentrated in the deepest three layers (10th, 11th, 12th). By comparing the distribution of social bias neurons for RoBERTa and FairBERTa, we observe that although there has been little change in quantity after debiasing (from 10.35 to 10.06 on average), there have been noteworthy alterations in the distribution of these neurons. Before debiasing, social bias neurons were predominantly concentrated in the deepest three layers (10th, 11th, 12th), with the highest concentration found in the deepest layer (12th), accounting for approximately 55.7% (BERT) and 36.4% (RoBERTa) of the total quantity. After debiasing, we notice that a substantial quantity (almost 1/3) of social bias neurons migrated from the deepest three layers (10th, 11th, 12th) to the shallowest three layers (1st, 2nd, 3rd). We observe a notable increase in the number of neurons in the two shallowest layers (1st, 2nd), rising from 2.9% (RoBERTa) to 26.0% (FairBERTa). Meanwhile, the number of neurons in the deepest layer (12th) decreased to 26.4%. Based on this phenomenon, we speculate that social bias neurons have a considerable adverse impact on the model outputs due to their proximity to the final output layer before debiasing. By pre-training on the anti-stereotypical data, FairBERTa transfers the social bias neurons from the deepest three layers to the shallowest three layers to mitigate social biases. We analyze that the significant reduction of social bias neurons near the final output layer alleviates their impacts on model outputs, which could be the secret lying behind the effectiveness of the debiasing method in FairBERTa.

5.2 STATISTICS OF SOCIAL BIAS NEURONS

Figure 4 presents the average number of social bias neurons pinpointed by our IG^2 for different demographic dimensions studied in this work. As shown, the average number varies for different demographic dimensions, with that for "Ethnicity" being the most (14.57) and that for "Gender"

Model	Ethnicity			Physical Appearance			Politics		
	Avg. BN	Avg. Intra	Avg. Inter	Avg. BN	Avg. Intra	Avg. Inter	Avg. BN	Avg. Intra	Avg. Inter
BERT	14.57	10.97	0.34	2.91	2.13	0.04	3.49	2.31	0.01
RoBERTa	14.05	8.01	0.38	3.17	1.13	0.02	4.75	3.06	0.00
FairBERTa	13.92	8.09	0.28	3.04	1.27	0.04	5.13	3.79	0.01

Table 6: Statistics of social bias neurons. "Avg. BN" means the average number of Bias Neurons pinpointed by our IG^2. "Avg. Intra" means the average number of neurons in the intersection within the same JUT (only prompt templates are different). "Avg. Inter" means the average number of neurons in the intersection for different JUTs (either different UTs or different Judgmental Modifiers).

being the fewest (1.09). Except for the "Ethnicity" dimension, we on average only need to suppress the activation of less than 4 neurons to effectively mitigate social biases for other dimensions.

Table 6 shows the statistics of social bias neurons. As we can see, for each demographic dimension, the same JUT shares most of the pinpointed social bias neurons, while different JUTs share almost no social bias neurons. The large value of Avg. Intra indicates that for the same JUT (confirmed Judgemental Modifier and Unfair Target), pinpointed social bias neurons remain almost the same for different prompt templates. This further verifies the accuracy of the social bias neurons pinpointed by our IG^2, because this suggests that pinpointed neurons are not related to the sentence structure or grammar of the prompt template, but only to the social bias in the studied social group. The small value of Avg. Inter indicates that social bias neurons pinpointed by our IG^2 are quite different for different JUTs, even if prompt templates are the same for them.
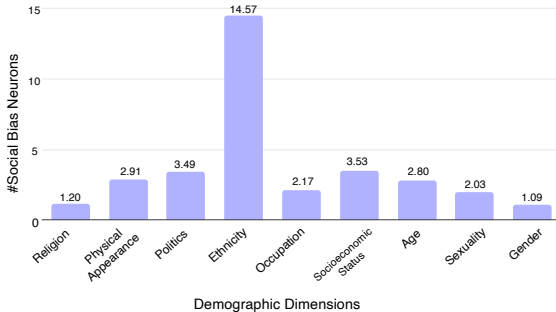

Figure 4: The average number of social bias neurons pinpointed in BERT for different demographic dimensions. Best viewed on screen.

6 RELATED WORK

Since various AI applications permeate every aspect of our lives (Zan et al., 2024; Liu et al., 2023b; Zan et al., 2023a; Chen et al., 2022; Liu et al., 2022; Zan et al., 2023b; Liu & Yang, 2021), research on AI Ethics (Liu et al., 2023c;d; Mehrabi et al., 2019) has attracted more and more attention. In this work, we mainly explore one important aspect of AI Ethics: AI Fairness, which has been studied from different perspectives (Hardt et al., 2016; John-Mathews et al., 2022; Nadeem et al., 2021; Nangia et al., 2020). (Liu et al., 2023a) proposed to study the existence of annotator group bias in various real-world crowdsourcing datasets. (Li et al., 2022) measured hierarchical regional bias in pre-trained language models. Some works tried to detect and mitigate social biases in word embeddings (Bolukbasi et al., 2016; Kaneko et al., 2022) and hidden representations (Chowdhury & Chaturvedi, 2022), while others explored quantifying social biases in downstream tasks. Many works have explored the fairness problem in text classification tasks (Dixon et al., 2018; Liu et al., 2021; Dinan et al., 2020). Some works also explore the fairness problem in generation tasks, such as machine translation (Stanovsky et al., 2019), story generation (Lucy & Bamman, 2021), and question answering (Parrish et al., 2022). However, no work has focused on the interpretability of the fairness research. In this paper, we close this gap by proposing an interpretable technique specific to the study of social bias along multiple dimensions.

7 LIMITATIONS AND CONCLUSION

In this paper, we propose a novel interpretable technique, Integrated Gap Gradients (IG^2), to precisely identify social bias neurons in pre-trained language models. We also develop a new dataset to facilitate the interpretability study of social bias. Derived from our interpretable technique, BIAS NEURON SUPPRESSION (BNS) is further proposed to mitigate social bias. Extensive experiments have verified the effectiveness of our IG^2 and BNS. In addition, facilitated by our interpretable

method, we analyze the distribution shift of social bias neurons after debiasing and obtain useful insights that bring inspiration to future fairness research.

Limitations. While our study provides valuable insights, we recognize there exist limitations. For example, our proposed BNS method directly sets the activation values of selected social bias neurons to zero. Although this is effective, designing a more refined suppression method might yield even better results. These present opportunities for future research.

REFERENCES

Afra Feyza Akyürek, Sejin Paik, Muhammed Yusuf Kocyigit, Seda Akbiyik, {S}erife Leman Runyun, and Derry Wijaya. On measuring social biases in prompt-based multi-task learning. 2022.

Haozhe An and Rachel Rudinger. Nichelle and nancy: The influence of demographic attributes and tokenization length on first name biases. May 2023.

Tolga Bolukbasi, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. Man is to computer programmer as woman is to homemaker? debiasing word embeddings, 2016. URL https://arxiv.org/abs/1607.06520.

Xiaokang Chen, Jiahui Chen, Yan Liu, and Gang Zeng. D^{3}etr: Decoder distillation for detection transformer, 2022.

Somnath Basu Roy Chowdhury and Snigdha Chaturvedi. Learning fair representations via rate-distortion maximization. *arXiv preprint arXiv:2202.00035*, 2022.

AidaMostafazadeh Davani, Mohammad Atari, Brendan Kennedy, and Morteza Dehghani. Hate speech classifiers learn human-like social stereotypes.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional Transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1*, pp. 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.

Emily Dinan, Angela Fan, Ledell Wu, Jason Weston, Douwe Kiela, and Adina Williams. Multi-dimensional gender bias classification. *arXiv preprint arXiv:2005.00614*, 2020.

Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 67–73, 2018.

Yanai Elazar, Nora Kassner, Shauli Ravfogel, Abhilasha Ravichander, Eduard H. Hovy, Hinrich Schütze, and Yoav Goldberg. Measuring and improving consistency in pretrained language models. *CoRR*, abs/2102.01017, 2021. URL https://arxiv.org/abs/2102.01017.

Shangbin Feng, ChanYoung Park, Yuhan Liu, and Yulia Tsvetkov. From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair nlp models. May 2023.

Eve Fleisig, Aubrie Amstutz, Chad Atalla, Su, Lin Blodgett, Hal Daumé, Alexandra Olteanu, Emily Sheng, Dan Vann, and Hanna Wallach. Fairprism: Evaluating fairness-related harms in text generation.

Seraphina Goldfarb-Tarrant, Adam Lopez, Roi Blanco, and Diego Marcheggiani. Bias beyond english: Counterfactual tests for bias in sentiment analysis in four languages. May 2023.

Yue Guo, Yi Yang, and Ahmed Abbasi. Auto-debias: Debiasing masked language models with automated biased prompts. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1012–1023, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.72. URL https://aclanthology.org/2022.acl-long.72.

Moritz Hardt, Eric Price, and Nathan Srebro. Equality of opportunity in supervised learning, 2016. URL https://arxiv.org/abs/1610.02413.

Minqing Hu and Bing Liu. Mining and summarizing customer reviews. In *KDD '04: Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 168–177, New York, NY, USA, 2004. ACM. ISBN 1-58113-888-1. doi: http://doi.acm.org/10.1145/1014052.1014073.

Jean-Marie John-Mathews, Dominique Cardon, and Christine Balagué. From reality to world. a critical perspective on ai fairness. *Journal of Business Ethics*, pp. 1–15, 2022.

Masahiro Kaneko and Danushka Bollegala. Debiasing pre-trained contextualised embeddings. *Cornell University - arXiv,Cornell University - arXiv*, Jan 2021.

Masahiro Kaneko, Danushka Bollegala, and Naoaki Okazaki. Gender bias in meta-embeddings. *arXiv preprint arXiv:2205.09867*, 2022.

Yizhi Li, Ge Zhang, Bohao Yang, Chenghua Lin, Anton Ragni, Shi Wang, and Jie Fu. HERB: Measuring hierarchical regional bias in pre-trained language models. In *Findings of the Association for Computational Linguistics: AACL-IJCNLP 2022*. Association for Computational Linguistics, November 2022. URL https://aclanthology.org/2022.findings-aacl.32.

Haochen Liu, Wei Jin, Hamid Karimi, Zitao Liu, and Jiliang Tang. The authors matter: Understanding and mitigating implicit bias in deep text classification. *arXiv preprint arXiv:2105.02778*, 2021.

Haochen Liu, Joseph Thekinen, Sinem Mollaoglu, Da Tang, Ji Yang, Youlong Cheng, Hui Liu, and Jiliang Tang. Toward annotator group bias in crowdsourcing. 2023a.

Yan Liu and Yazheng Yang. Enhance long text understanding via distilled gist detector from abstractive summarization, 2021.

Yan Liu, Sanyuan Chen, Yazheng Yang, and Qi Dai. MPII: Multi-level mutual promotion for inference and interpretation. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7074–7084, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-long.488. URL https://aclanthology.org/2022.acl-long.488.

Yan Liu, Xiaokang Chen, and Qi Dai. Parallel sentence-level explanation generation for real-world low-resource scenarios. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023b. doi: 10.1109/ICASSP49357.2023.10094616.

Yan Liu, Xiaokang Chen, Yan Gao, Zhe Su, Fengji Zhang, Daoguang Zan, Jian-Guang Lou, Pin-Yu Chen, and Tsung-Yi Ho. Uncovering and quantifying social biases in code generation, 2023c.

Yan Liu, Yan Gao, Zhe Su, Xiaokang Chen, Elliott Ash, and Jian-Guang Lou. Uncovering and categorizing social biases in text-to-sql, 2023d.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A robustly optimized BERT pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

Li Lucy and David Bamman. Gender and representation bias in GPT-3 generated stories. In *Proceedings of the Third Workshop on Narrative Understanding*, pp. 48–55, Virtual, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.nuse-1.5. URL https://aclanthology.org/2021.nuse-1.5.

Nicholas Meade, Elinor Poole-Dayan, and Siva Reddy. An empirical survey of the effectiveness of debiasing techniques for pre-trained language models, 2021. URL https://arxiv.org/abs/2110.08527.

Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *arXiv: Learning*, 2019.

Moin Nadeem, Anna Bethke, and Siva Reddy. Stereoset: Measuring stereotypical bias in pretrained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Jan 2021. doi: 10.18653/v1/2021.acl-long.416. URL `http://dx.doi.org/10.18653/v1/2021.acl-long.416`.

Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. CrowS-pairs: A challenge dataset for measuring social biases in masked language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.154. URL `https://aclanthology.org/2020.emnlp-main.154`.

Nedjma Ousidhoum, Xinran Zhao, Tianqing Fang, Yangqiu Song, and Dit-Yan Yeung. Probing toxic content in large pre-trained language models. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, Online, August 2021. Association for Computational Linguistics.

Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel Bowman. BBQ: A hand-built bias benchmark for question answering. In *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 2086–2105, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-acl.165. URL `https://aclanthology.org/2022.findings-acl.165`.

Rebecca Qian, Candace Ross, Jude Fernandes, Eric Smith, Douwe Kiela, and Adina Williams. Perturbation augmentation for fairer nlp. May 2022.

Omar Shaikh, Hongxin Zhang, William Held, Michael Bernstein, and Diyi Yang. On second thought, let's not think step by step! bias and toxicity in zero-shot reasoning. Dec 2022.

Emily Sheng, Kai-Wei Chang, Premkumar Natarajan, and Nanyun Peng. The woman worked as a babysitter: On biases in language generation. *empirical methods in natural language processing*, 2019.

Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. Evaluating gender bias in machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1679–1684, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1164. URL `https://aclanthology.org/P19-1164`.

Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3319–3328. PMLR, 2017. URL `http://proceedings.mlr.press/v70/sundararajan17a.html`.

Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart Shieber. Investigating gender bias in language models using causal mediation analysis. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 12388–12401. Curran Associates, Inc., 2020. URL `https://proceedings.neurips.cc/paper/2020/file/92650b2e92217715fe312e6fa7b90d82-Paper.pdf`.

Kellie Webster, Xuezhi Wang, Ian Tenney, Alex Beutel, Emily Pitler, Ellie Pavlick, Jilin Chen, and Slav Petrov. Measuring and reducing gendered correlations in pre-trained models. *arXiv: Computation and Language*, 2020.

Daoguang Zan, Bei Chen, Yongshun Gong, Junzhi Cao, Fengji Zhang, Bingchao Wu, Bei Guan, Yilong Yin, and Yongji Wang. Private-library-oriented code generation with large language models, 2023a.

Daoguang Zan, Ailun Yu, Bo Shen, Jiaxin Zhang, Taihong Chen, Bing Geng, Bei Chen, Jichuan Ji, Yafen Yao, Yongji Wang, and Qianxiang Wang. Can programming languages boost each other via instruction tuning?, 2023b.

Daoguang Zan, Ailun Yu, Wei Liu, Dong Chen, Bo Shen, Wei Li, Yafen Yao, Yongshun Gong, Xiaolin Chen, Bei Guan, Zhiguang Yang, Yongji Wang, Qianxiang Wang, and Lizhen Cui. Codes: Natural language to code repository via multi-layer sketch, 2024.