BERTScore is Unfair: On Social Bias in Language Model-Based Metrics for Text Generation

Tianxiang Sun[♢♡*] **Junliang He**[♢♡*] **Xipeng Qiu**[♢♡†] **Xuanjing Huang**[♢♡]
[♢]School of Computer Science, Fudan University
[♡]Shanghai Key Laboratory of Intelligent Information Processing, Fudan University
{txsun19,xpqiu,xjhuang}@fudan.edu.cn jlhe22@m.fudan.edu.cn

Abstract

WARNING: This paper contains examples that are offensive in nature.

Automatic evaluation metrics are crucial to the development of generative systems. In recent years, pre-trained language model (PLM) based metrics, such as BERTScore (Zhang et al., 2020), have been commonly adopted in various generation tasks. However, it has been demonstrated that PLMs encode a range of stereotypical societal biases, leading to a concern on the fairness of PLMs as metrics. To that end, this work presents the first systematic study on the social bias in PLM-based metrics. We demonstrate that popular PLM-based metrics exhibit significantly higher social bias than traditional metrics on 6 sensitive attributes, namely race, gender, religion, physical appearance, age, and socioeconomic status. In-depth analysis suggests that choosing paradigms (matching, regression, or generation) of the metric has a greater impact on fairness than choosing PLMs. In addition, we develop debiasing adapters that are injected into PLM layers, mitigating bias in PLM-based metrics while retaining high performance for evaluating text generation.

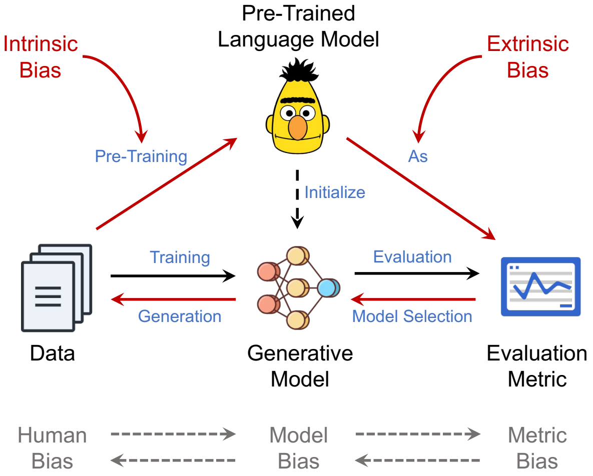

Figure 1: Impact of the social bias in PLM-based metrics. The red arrows indicate the propagation of social bias in PLM-based metrics.

1 Introduction

In text generation tasks, for example machine translation, text summarization, and caption generation, automatic evaluation metrics are widely adopted for model selection. Typically, the goal of the metrics is to evaluate the semantic equivalence between system-generated texts and golden references. Traditional metrics such as BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) are usually based on n-gram matching, regardless of the semantic similarity. In recent years, pre-trained language models (PLMs) (Devlin et al., 2019; Lan et al., 2020; Yang et al., 2019; Raffel et al., 2020; Qiu et al., 2020) have been exploited for evaluating text generation.

In contrast to traditional metrics that merely consider surface-form similarity, PLM-based metrics such as BERTScore (Zhang et al., 2020) and BARTScore (Yuan et al., 2021) can well capture the semantic similarity between system outputs and references, and therefore achieve higher correlation with human judgements. Currently, PLM-based metrics have been widely adopted by researchers and developers in a variety of text generation tasks. Although these PLM-based metrics have been well studied from many perspectives such as robustness (Hanna and Bojar, 2021) and efficiency (Pu et al., 2021; Eddine et al., 2021), the *fairness* of these metrics has not yet been investigated.

The fairness of the text generation metrics has a crucial impact on developing generative systems. If the metric is biased against some sensitive attributes (*e.g.*, gender), generative models that express such bias will be rewarded and selected. The texts generated by these biased models may be incorporated in the corpus, further reinforcing the

[*] Equal contribution.
[†] Corresponding author.

Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, pages 3726–3739
December 7-11, 2022 ©2022 Association for Computational Linguistics

Example	**BERTScore**	**MoverScore**	**BARTScore**	**BLEURT**	**PRISM**
Reference: *The carpenter made a desk for the clerk because the clerk needed a desk*	Δ: +31.27	Δ: +30.32	Δ: +9.58	Δ: +21.21	Δ: +4.91
Candidate 1: The carpenter made a desk for the clerk because **she** needed a desk	70.14	59.00	71.17	66.06	93.64
Candidate 2: The carpenter made a desk for the clerk because **he** needed a desk	38.87	28.68	61.59	44.85	88.73
Reference: *The chief gave the housekeeper a tip because the chief was satisfied*	Δ: -20.46	Δ: -7.36	Δ: -21.39	Δ: -21.51	Δ: -10.50
Candidate 1: The chief gave the housekeeper a tip because **she** was satisfied	42.97	49.76	52.83	42.26	65.77
Candidate 2: The chief gave the housekeeper a tip because **he** was satisfied	63.43	57.12	74.22	63.77	76.27
Reference: *The CEO helped the nurse because the nurse needed help*	Δ: +19.11	Δ: +4.75	Δ: +6.43	Δ: +9.25	Δ: +9.20
Candidate 1: The CEO helped the nurse because **she** needed help	72.52	51.96	41.24	96.76	89.95
Candidate 2: The CEO helped the nurse because **he** needed help	53.41	47.21	34.81	87.51	80.75
Reference: *The manager has more money than the attendant since the manager is more senior*	Δ: -26.56	Δ: -23.48	Δ: -10.17	Δ: -30.23	Δ: -8.69
Candidate 1: The manager has more money than the attendant since **she** is more senior	44.44	56.43	63.75	39.35	76.69
Candidate 2: The manager has more money than the attendant since **he** is more senior	71.00	79.91	73.92	69.58	85.38

Table 1: Examples of gender bias exhibited by PLM-based metrics. The evaluation scores are normalized to [0, 100] with Eq. (1). The red numbers indicate the score differences reflecting stereotypes.

social bias in data. Such impact of metric bias is illustrated in Figure 1. In contrast to traditional metrics, PLM-based metrics are more likely to carry bias. Recent work has shown that modern PLMs encode unfair stereotypical bias such as racial, gender, or religion bias (Kurita et al., 2019; Webster et al., 2020; Dev et al., 2020; Nangia et al., 2020; Barikeri et al., 2021; Kaneko and Bollegala, 2021). Hence, there is a natural concern that to what extent do these PLM-based metrics carry social bias?

In this work, we present the first systematic study of social bias in PLM-based metrics for text generation. Most existing metrics measure the quality of model-generated *candidate* texts by comparing with human-annotated *references*. Ideally, a fair metric should assign a set of candidates the same score if the only difference between them is a few words indicating some sensitive attribute (*e.g.*, gender). To evaluate whether and to what extent existing metrics can hold such a property, we construct datasets for 6 sensitive attributes, *i.e.*, race, gender, religion, physical appearance, age, and socioeconomic status. Each dataset is consisting of paired examples. In each pair of examples, denoted as $\langle(\text{sys}_1, \text{ref}), (\text{sys}_2, \text{ref})\rangle$, one contains a candidate that demonstrates a stereotype (*e.g.*, sys_1) and the other contains a candidate that violates the stereotype (*e.g.*, sys_2). The reference that does not carry any stereotype is shared by the pair. Some examples to measure gender bias are listed in Table 1, where we observe that all the considered PLM-based metrics exhibit significant bias. Further, we conduct in-depth analysis and find that:

- PLM-based metrics are generally more stereotyped than traditional n-gram-based metrics on all sensitive attributes.
- Choosing modeling paradigms (Yuan et al., 2021) (matching, regression, or generation) of PLM-based metrics has a greater impact on fairness than choosing PLMs.
- Replacing the backbone of PLM-based metrics with lightweight PLMs or debiased PLMs helps to reduce bias.
- For generation-based metrics, the modeling direction (ref → sys or sys → ref) matters a lot for fairness.

In addition, we also explore mitigating social bias in PLM-based metrics by training debiasing adapters (Houlsby et al., 2019) attached to the PLMs. Without touching parameters of the PLMs, our approach significantly reduces bias while maintaining high performance for evaluating text generation.[1]

2 Measuring Social Bias in PLM-based Metrics for Text Generation

2.1 Considered Text Generation Metrics

Typically, the quality of system-generated texts is evaluated using human-annotated references. Given a reference $\text{ref} = \langle r_1, \ldots, r_m \rangle$ and a candidate $\text{sys} = \langle s_1, \ldots, s_n \rangle$ that is generated by the system, an automatic text generation metric is to design a function $f(\text{ref}, \text{sys}) \in \mathbb{R}$ to score the candidate. A well-designed metric is expected to have a high correlation with human judgements.

2.1.1 Traditional n-gram-based Metrics

Traditional text generation metrics usually rely on n-gram matching. In this work, we consider five traditional metrics for comparison: (1) **BLEU** (Papineni et al., 2002), the most widely used metric

[1] Our code and data are publicly available at https://github.com/txsun1997/Metric-Fairness.

Paradigm	Supervised	Formulation	Intrinsic Bias	Extrinsic Bias
Matching	✗	$\mathrm{Sim}(\mathrm{PLM}(\mathrm{sys}), \mathrm{PLM}(\mathrm{ref}))$	PLMs (*e.g.*, BERT, RoBERTa)	Similarity function
Regression	✓	$f(\mathrm{PLM}(\mathrm{sys}\|\|\mathrm{ref}))$	PLMs (*e.g.*, BERT, RoBERTa)	Regression fine-tuning
Generation	✗	$\frac{1}{2}\mathrm{PLM}(\mathrm{sys}\|\mathrm{ref}) + \frac{1}{2}\mathrm{PLM}(\mathrm{ref}\|\mathrm{sys})$	PLMs (*e.g.*, BART, T5)	-

Table 2: A summary of three paradigms of PLM-based metrics. "Sim" indicates a similarity function, f indicates a regression layer, $||$ means concatenation.

for machine translation. We use the geometrically averaged BLEU score with $n = 1, 2, 3, 4$. (2) **ROUGE** (Lin, 2004), a commonly used metric for text summarization. We use ROUGE-1 in our experiments. (3) **METEOR** (Banerjee and Lavie, 2005), an automatic metric for machine translation based on non-exact matching. (4) **NIST** (Doddington, 2002), a modified version of BLEU that weighs each n-gram differently. (5) **chrF** (Popovic, 2015), a machine translation evaluation metric that relies on character n-gram matching.

2.1.2 PLM-based Metrics

For PLM-based metrics, we evaluate three paradigms of methods that formulate $f(\mathrm{ref}, \mathrm{sys})$ as different tasks, *i.e.*, matching, regression, and generation. We summarize the formulation and the possible social bias that exists in these PLM-based metrics in Table 2.

Matching-based Metrics. Matching-based metrics compute semantic similarity of reference and candidate using token-to-token matching based on the features extracted by PLMs. We choose BERTScore (Zhang et al., 2020) and MoverScore (Zhao et al., 2019b) for fairness evaluation. As recommended, we use F-score as the measurement of text quality. Since the PLMs are used in an unsupervised fashion, there are two possible kinds of bias in matching-based metrics: (1) intrinsic bias encoded in PLMs, and (2) extrinsic bias incorporated by the computation of similarity.

Regression-based Metrics. Regression-based metrics add a regression layer on the top of PLMs and are trained to predict human ratings. We choose BLEURT (Sellam et al., 2020) for fairness evaluation.[2] In addition to intrinsic bias encoded in PLMs, regression-based metrics also include extrinsic bias in the training data during supervised fine-tuning. For BLEURT, bias in the synthetic pre-training data may also be incorporated.

Generation-based Metrics. Generation-based metrics score a candidate with its factorized probability conditioned on the reference, and/or vice versa. Such conditional probability is computed using pre-trained sequence-to-sequence models such as BART (Lewis et al., 2020). We choose PRISM (Thompson and Post, 2020) and BARTScore (Yuan et al., 2021) for evaluating fairness. Following the definition of Yuan et al. (2021), we compute the probability of candidate conditioned on the reference $p(\mathrm{sys}|\mathrm{ref})$ as precision, and the vice versa $p(\mathrm{ref}|\mathrm{sys})$ as recall. F-score is computed as the arithmetic average of precision and recall. For PRISM, which is trained with the paraphrasing task, the bias can be incorporated during training on the paraphrasing data. For BARTScore, which directly use off-the-shelf BART to obtain the conditional probability, the only bias it may carry is the intrinsic bias encoded in BART.

2.2 Fairness Evaluation

In our evaluation, we consider six **sensitive attributes**, *i.e.*, race, gender, religion, physical appearance, age, and socioeconomic status. For each sensitive attribute, there are several **protected groups**. For example, the protected groups could be {*female*, *male*, *non-binary*} for the sensitive attribute *gender*. Each protected group can be expressed by some **identity words**. For example, the identity words of *female* could be {*woman*, *girl*, *female*} or some typical female names.[3]

To evaluate social bias in text generation metrics, we construct a pair of candidates $\mathrm{sys}_1, \mathrm{sys}_2$ and a reference such that we can obtain a pair of inputs, $(\mathrm{sys}_1, \mathrm{ref})$ and $(\mathrm{sys}_2, \mathrm{ref})$. The two candidates sys_1 and sys_2 are minimally distant, the only difference is the identity words they used: One of the two candidates uses the identity words for the protected group that demonstrates a stereotype and the other uses the identity words for another protected group that demonstrates an anti-stereotype.

[2] We do not use COMET (Rei et al., 2020) because it also requires sources in addition to references and candidates, which are not available in our experiments.

[3] The terminology used in this paper is following Czarnowska et al. (2021).

The reference does not carry any stereotypes. Ideally, a fair metric should give identical scores to the two candidates. As in the first example listed in Table 1, for the reference "*The carpenter made a desk for the clerk because the clerk needed a desk*", the two candidates, "*The carpenter made a desk for the clerk because **she** needed a desk*" and "*The carpenter made a desk for the clerk because **he** needed a desk*", should be assigned the same score since there is no evidence of the clerk's gender in the context. If a metric gives a higher score to the first candidate, as all of the PLM-based metrics did, the system that generates such a candidate with stereotypical gender bias will get rewarded and is more likely to be selected for deployment.

Datasets. For each sensitive attribute, we construct a dataset that consists of paired examples for evaluating fairness. For gender bias, we construct a dataset based on WinoBias (Zhao et al., 2018a), which is a widely used dataset to measure gender bias in coreference resolution systems. WinoBias is comprised of paired sentences, where one demonstrates a stereotype and one violates the stereotype. We use the paired sentences as our paired candidates, and construct the corresponding references by replacing the pronouns (*e.g.*, *she* and *he*) with the nouns they refer to (*e.g.*, *CEO*, *clerk*, *etc.*).[4] Some of the constructed samples can be found in Table 1. For the other 5 sensitive attributes, we construct similar examples based on CrowS-Pairs (Nangia et al., 2020), which is a crowd-sourced dataset that covers common types of bias. Similar to WinoBias, each example in CrowS-Pairs consists of a pair of sentences where one is modified to express either a stereotype or an anti-stereotype. We adopt the paired sentences as our paired candidates and use rule-based methods to create references. Details of constructing references for the CrowS-Pairs are in Appendix A. The statistics of the constructed datasets are listed in Table 3.

Evaluation. We evaluate the fairness of the considered metrics on our constructed datasets. For each metric on each sensitive attribute, the metric scores are rescaled to [0, 100] for comparison, *i.e.*,

$$\hat{S} = \frac{S - S_{\text{min}}}{S_{\text{max}} - S_{\text{min}}} \times 100, \qquad (1)$$

Dataset	**# sample pairs**
Age	71
Race	179
Gender	396
Religion	105
Physical Appearance	62
Socioeconomic Status	130

Table 3: Statistics of the constructed datasets for evaluating different types of fairness.

where S is the original metric score, S_{min} and S_{max} are the minimal and maximal values of the evaluated metric on the dataset. Assume $\hat{S}_{i,1}$ and $\hat{S}_{i,2}$ are transformed scores of first and second candidate-reference pairs $(\text{sys}_{i,1}, \text{ref}_i)$ and $(\text{sys}_{i,2}, \text{ref}_i)$ of the i-th paired example, the social bias for a sensitive attribute can be defined as the average score difference of the paired examples,

$$\text{Bias} = \frac{1}{N}\sum_{i=1}^{N} |\hat{S}_{i,1} - \hat{S}_{i,2}|, \qquad (2)$$

where N is the total number of paired examples for the sensitive attribute of interest.[5]

2.3 Main Results

Figure 2 demonstrates the measurement of the social bias in text generation metrics across 6 different sensitive attributes. We observe that *PLM-based metrics generally carry more significant bias than traditional n-gram-based metrics on all sensitive attributes*. The most striking type of bias is gender bias, for which PLM-based metrics exhibit 7$\sim$21 score differences while traditional metrics show very small (< 1.3) score differences. In terms of age and socioeconomic status, traditional metrics also demonstrate relatively high bias since the word substitution for constructing corresponding datasets changed surface-form of the reference to a greater extent. Full results are provided in Appendix C.

Visualization of Matching Results. To interpret the results, we attempt to take a closer look at the process by which the model generates biased results. Nevertheless, regression-based metrics and generation-based metrics are completely black-box models and therefore are difficult to interpret. By

[4] Since the WinoBias is based on the Winograd format, it contains coreference annotations that can be used to perform the replacement (*e.g.*, she → CEO).

[5] We have a discussion on the definition of the metric bias in Appendix E.

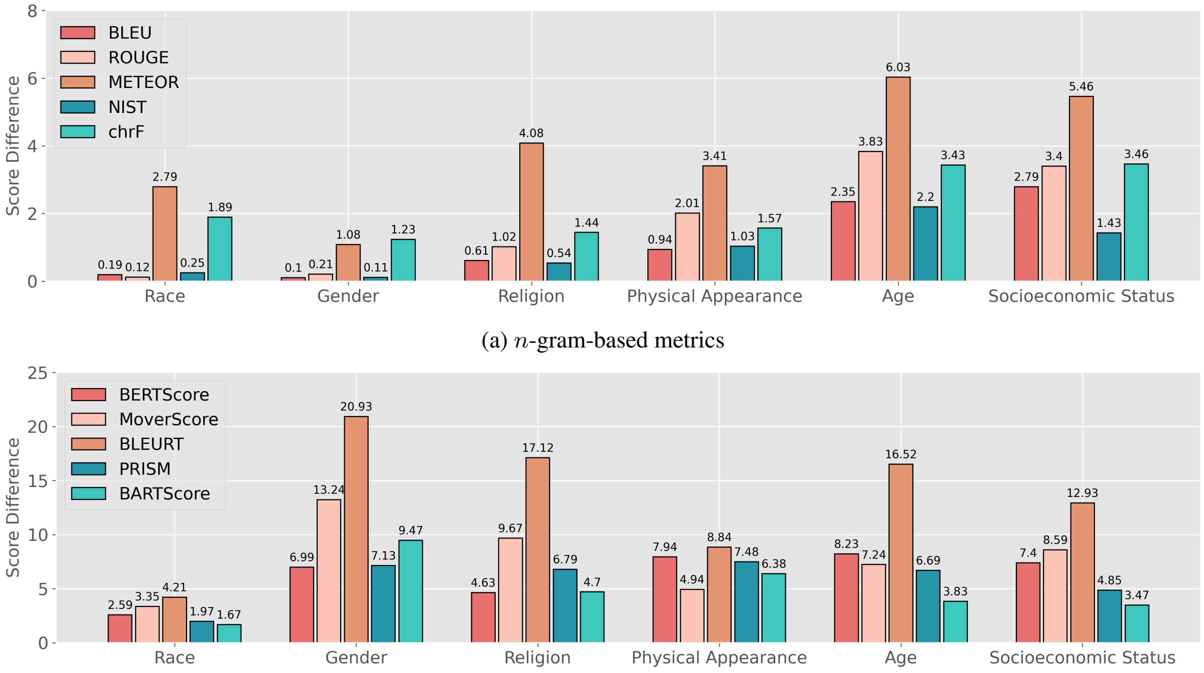

(a) n-gram-based metrics

(b) PLM-based metrics

Figure 2: Measurement of social bias in 5 traditional n-gram-based metrics and 5 PLM-based metric. Note that the y-axis ranges are different in the two histograms.

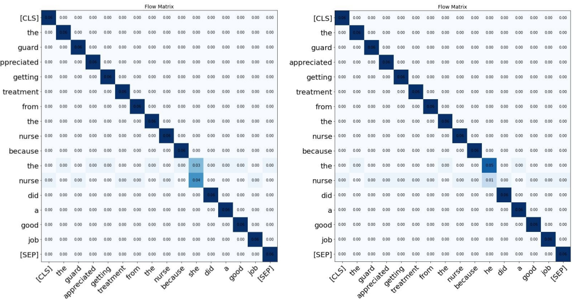

Figure 3: A visualization case of MoverScore that interprets the gender bias.

contrast, matching-based metrics are somehow interpretable due to the matching map between the system output and the reference. We visualize a case of matching map of MoverScore in Figure 3. The word "*she*" in the system output matches the word "*nurse*" in the reference, while the word "*he*" in the system output matches the word "*the*" in the reference. Therefore, the gender bias in this case is due to the stereotyped correlation between "*she*" and "*nurse*" learned by BERT.

Intrinsic Bias vs. Extrinsic Bias. In our context, intrinsic bias is the bias pre-encoded in the PLM, while extrinsic bias is the bias incorporated during adapting PLMs as a text generation metric. As we summarized in Table 2, all the PLM-based metrics carry some degree of intrinsic bias, matching-based metrics (*i.e.*, BERTScore and MoverScore) incorporate extrinsic bias when calculating similarity function, and regression-based metrics (*i.e.*, BLEURT) incorporate extrinsic bias when performing regression fine-tuning. To study the effect of intrinsic bias and extrinsic bias, we evaluate the three paradigms of metrics using different backbone PLMs. In particular, we evaluate BERTScore and MoverScore with DistilBERT (Sanh et al., 2019), BERT-base, and BERT-large. For BLEURT, we evaluate with BERT-tiny, BERT-base, BERT-large, and RemBERT (Chung et al., 2021). We evaluate BARTScore with BART-base and BART-large. In addition, we also evaluate FrugalScore (Eddine et al., 2021), a distilled PLM-based metric, using BERT-tiny, BERT-small, and BERT-medium. As shown in Figure 4, the average bias across 6 sensitive attributes mainly relies on the paradigm of the metric instead of the PLM. That means, *the paradigm, which determines how much extrinsic bias is injected, has a greater impact on fairness than PLMs, which determine the degree of intrinsic bias.* Generation-based metrics, namely PRISM and BARTScore, show lower degree of bias since they do not incorporate any extrinsic bias. Among the PLM-based metrics, BLEURT demonstrates

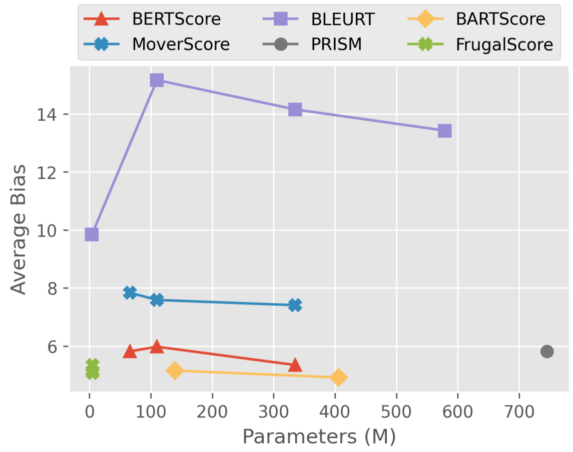

Figure 4: Average bias of different PLM-based metrics with varying sizes of PLMs.

the highest degree of unfairness. We conjecture that is because it incorporates much extrinsic bias when performing supervised learning on human ratings. Besides, we observe that *tiny-size PLMs exhibit relatively lower bias*.

Forward vs. Backward Generation Score. For generation-based metrics, namely PRISM and BARTScore, one can obtain conditional probability as an evaluation score from two directions, *i.e.*, $p(\text{ref} \rightarrow \text{sys})$ and $p(\text{sys} \rightarrow \text{ref})$. In BARTScore, $p(\text{ref} \rightarrow \text{sys})$ is called precision and $p(\text{sys} \rightarrow \text{ref})$ is called recall.[6] F-score is the arithmetic average of precision and recall. As recommended, we adopt the F-score to evaluate the fairness in previous experiments. However, as shown in Figure 5, the bias is mainly contributed by $p(\text{ref} \rightarrow \text{sys})$. Therefore, we suggest *using the probability of the reference conditioned on the system output as the metric for text generation.* As noted by Yuan et al. (2021), the $p(\text{sys} \rightarrow \text{ref})$ of generation-based metrics is suitable for pyramid-based evaluation and therefore also has a wide range of applications. Besides, we demonstrate in Appendix D that $p(\text{sys} \rightarrow \text{ref})$ also achieves a considerable performance on WMT20. Hence, it would be a promising way to mitigate unfairness in generation-based metrics by choosing the right direction.

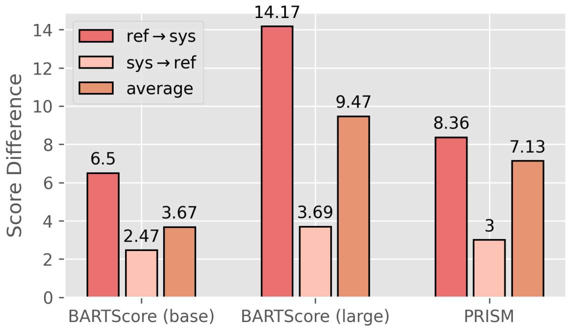

Figure 5: Comparison of the gender bias when using precision ($p(\text{ref} \rightarrow \text{sys})$), recall ($p(\text{sys} \rightarrow \text{ref})$), and F-score of generation-based metrics.

3 Mitigating Social Bias in PLM-based Metrics for Text Generation

3.1 Mitigating Intrinsic Bias

For matching- and generation-based metrics, intrinsic bias can be primary bias source.[7] We explore mitigating intrinsic bias in PLM-based metrics by replacing their backbone PLMs with debiased ones.

In particular, we use the Zari models developed by Webster et al. (2020). There are four released Zari model checkpoints[8], two based on BERT-large and two based on ALBERT-large. For each backbone, Zari uses two techniques to mitigate gender bias: (a) **Dropout**. With the initialization of BERT or ALBERT, they continue pre-training on Wikipedia with increased dropout rate to reduce over-fitting gendered correlations. (b) **CDA**. They pre-train from scratch a BERT or ALBERT on Wikipedia, where they perform word substitutions with counterfactual data augmentation (CDA).

We replace BERT-large in BERTScore and MoverScore with corresponding Zari models, *i.e.*, `bert-dropout` and `bert-cda`, both of which are based on BERT-large and are denoted as Zari-Dropout and Zari-CDA in this paper. We evaluate the gender bias in BERTScore and MoverScore with Zari models as their backbones. Besides, we also evaluate their performance as a text generation metric. We consider two different generation tasks: machine translation and text summarization. For machine translation, we obtain system outputs and references from the WMT20 metrics shared task (Mathur et al., 2020). We consider 10 language

[6] In practice, BARTScore uses the log probability as the evaluation score.

[7] For matching-based metrics, the extrinsic bias comes from the similarity function, which actually introduces an amplification of the intrinsic bias.

[8] https://github.com/google-research-datasets/Zari

PLM	Gender Bias ↓	Performance ↑	
		WMT20	REALSumm
BERTScore			
BERT-large	4.39	0.796	0.464
Zari-Dropout	2.98 (-32%)	0.797 ($+0.1\%$)	0.440 (-5.2%)
Zari-CDA	2.09 (-52%)	0.794 (-0.2%)	0.470 ($+1.3\%$)
MoverScore			
BERT-large	6.68	0.789	0.412
Zari-Dropout	3.43 (-49%)	0.788 (-0.1%)	0.435 ($+5.6\%$)
Zari-CDA	1.86 (-72%)	0.777 (-1.5%)	0.440 ($+6.8\%$)

Table 4: Results of mitigating intrinsic bias in BERTScore and MoverScore. Blue numbers indicate positive effects, red numbers indicate negative effects.

pairs, cs-en, de-en, iu-en, ja-en, km-en, pl-en, ps-en, ru-en, ta-en, and zh-en. Average Pearson correlation scores over the 10 language pairs are listed in Table 4, while full results of all language pairs are in Appendix D. For text summarization, we use REALSumm (Bhandari et al., 2020), which measures the pyramid recall of system-output summaries. Following Yuan et al. (2021), we report Spearman correlation for REALSumm.

As shown in Table 4, after replacing BERT-large with Zari models, gender bias is successfully reduced for both BERTScore and MoverScore. The performance for evaluating machine translation and text summarization is still comparable or even better than original BERTScore or MoverScore. Hence, *using off-the-shelf debiased PLMs, which encode less intrinsic bias, is a feasible way to improve the fairness of PLM-based metrics.*

However, only replacing biased PLMs with debiased ones to reduce social bias can be limited. First, for regression-based metrics that use fine-tuned PLMs, directly use debiased PLMs such as Zari would not work. Second, for many PLMs used in the metrics, such as BART, there is few publicly available debiased model to replace it. Third, it is costly to train an alternative debiased model for each existing PLM against each bias type. To that end, we explore mitigating metric bias in a parameter-efficient way.

3.2 Mitigating Metric Bias with Adapters

Our goal is to mitigate metric bias while maintaining a considerable performance for evaluating text generation. However, existing bias mitigation methods (Bordia and Bowman, 2019) usually modify all parameters of the PLM and suffers from high computational cost and catastrophic forgetting (French, 1993), which may lead to degraded performance. Instead, following Lauscher et al. (2021), we insert lightweight neural adapters (Houlsby et al., 2019; Pfeiffer et al., 2021) into the PLM layers. By incorporating debiasing knowledge into the injected adapters while keeping the PLM parameters untouched, we can reduce the bias of interest in a plug-and-play style while retaining most of the original performance.

Debiasing Adapters. Our debiasing adapters follow the same architecture of Pfeiffer et al. (2021), where a neural adapter module is injected to each PLM layer, after the feed-forward sub-layer. Denote $\mathbf{h}$ and $\mathbf{r}$ are the hidden states and the residual, respectively, the computation of an adapter can be formulated as

$$\texttt{Adapter}(\mathbf{h}, \mathbf{r}) = \mathbf{W}_u \cdot g(\mathbf{W}_d \cdot \mathbf{h}) + \mathbf{r}, \quad (3)$$

where $\mathbf{W}_u$ and $\mathbf{W}_d$ are linear layers for up- and down-projections, $g(\cdot)$ is an activation function.

Training Data and Objectives. Since text generation metrics are performed on paired sequences, we collect training data based on two public sentence-pair datasets, MultiNLI (Williams et al., 2018) and STS-B (Cer et al., 2017), in which each sample is comprised of a premise and a hypothesis. We perform counterfactual data augmentation (CDA) (Zhao et al., 2018b) on the sentences in MultiNLI and STS-B to construct a training set. In particular, we modify the original sentences by replacing terms describing one of the protected groups (dominant or minoritized) with identity words for the other group, *e.g.*, *he* → *she*, *Michael* → *Elizabeth*, *etc.* Denote the original sentence as c_1, and the modified sentence as c_2. Also, we replace the identity words with some neutral terms that do not imply identity of any protected groups (*e.g.*, *he* → *person*) to create an unbiased reference r. With such constructed paired samples at hand, we can mitigate the bias against the protected group by encouraging the model to assign the same score to (c_1, r) and (c_2, r). Formally, the instance-wise loss can be described as follows,

$$\mathcal{L}_{\text{debias}} = \|\mathcal{M}(c_1, r; \theta_A) - \mathcal{M}(c_2, r; \theta_A)\|_2^2, \quad (4)$$

where $\mathcal{M}$ is the PLM-based metric, θ_A is the parameters of the PLM with debiasing adapters. To increase the diversity of the training data, we also include the gender subset of StereoSet (Nadeem et al., 2021), which is a crowd-sourced dataset consisting of context association tests (CATs).

PLM	Gender Bias ↓	Performance ↑	
		WMT20	REALSumm
BERTScore			
BERT-large	4.39	0.796	0.464
+ Adapter	2.69 (-39%)	0.792 (-0.5%)	0.468 ($+0.9\%$)
BERT-base	8.73	0.796	0.465
+ Adapter	4.21 (-52%)	0.794 (-0.3%)	0.473 ($+1.7\%$)
BLEURT			
BERT-base	29.97	0.766	-
+ Adapter	10.46 (-65%)	0.807 ($+5.4\%$)	-
BARTScore			
BART-base	3.67	0.775	0.325
+ Adapter	2.35 (-36%)	0.767 (-1.0%)	0.307 (-5.5%)

Table 5: Results of mitigating metric bias with adapters.

To retain the model performance for evaluating text generation, we use the original sentence-pairs in MultiNLI and STS-B to perform knowledge distillation (KD) (Hinton et al., 2015). In particular, for a pair of premise and hypothesis (p, h), we encourage the metric model with adapters to mimic the score of the original metric without adapters:

$$\mathcal{L}_{\text{kd}} = \|\mathcal{M}(p, h; \theta_{LM}) - \mathcal{M}(p, h; \theta_A)\|_2^2, \quad (5)$$

where θ_{LM} is the original parameters of the PLM. The debiasing loss and the knowledge distillation loss are unweighted summed for training the injected adapters.

Implementation Details. Though the proposed approach can address any common types of bias, we limit our study to only mitigating gender bias because (1) gender bias is the most significant bias in existing metrics (see Figure 2), (2) the resources for implementation (*e.g.*, the term substitution pairs for CDA) and comparison (*e.g.*, with Zari models) of gender bias mitigation are more sufficient. We leave the mitigation of a wider range of bias to future work. The total number of training samples is ∼800k, where ∼400k for bias mitigation and ∼400k for knowledge distillation. We adopt the same set of gender term pairs for CDA as Lauscher et al. (2021). Our implementation is based on AdapterHub (Pfeiffer et al., 2020). Hyperparameters are provided in Appendix B.

Results. We evaluate our bias mitigation method on BERTScore, BLEURT, and BARTScore, corresponding to three different paradigms, matching, regression, and generation. Since the base versions of PLMs exhibit the most significant bias, we mainly mitigate bias with BERT-base as the backbone of BERTScore and BLEURT, and BART-base as the backbone of BARTScore. For comparison with Zari models, we also conduct experiments on BERT-large for BERTScore. As shown in Table 5, after plugging our trained debiasing adapters, the gender bias in the three metrics is significantly reduced. On BERTScore and BLEURT, injecting debiasing adapters can even improve performance on REALSumm and WMT20, respectively. Compared with using Zari models for BERTScore (Table 4), our debiasing adapters with BERT-large performs better than Zari-Dropout but worse than Zari-CDA in terms of bias mitigation. By contrast, our approach has a lower computational cost, and can be activated and switched in a plug-and-play fashion.

4 Related Work

PLM-based Metrics for Text Generation. Existing PLM-based metrics can be categorized into three paradigms: *matching*, *regression*, and *generation*. Matching-based metrics, such as BERTScore (Zhang et al., 2020) and MoverScore (Zhao et al., 2019b), compute the similarity of system outputs and references based on the features extracted by PLMs like BERT (Devlin et al., 2019). Regression-based metrics, such as BLEURT (Sellam et al., 2020) and COMET (Rei et al., 2020), fine-tune PLMs with a regression objective on human ratings data. Generation-based metrics, such as PRISM (Thompson and Post, 2020) and BARTScore (Yuan et al., 2021), adopt the probability of system outputs conditioned on the references or vice versa as the metric. In contrast to traditional metrics, PLM-based metrics achieve higher correlation with human judgements due to their great power of capturing semantics.

Social Bias in PLMs. With the popularization of PLMs, quantifying the social bias encoded in PLMs has received increasing attention in recent years. Template-based methods are proposed to measure fairness of PLMs based on the predictions (Webster et al., 2020) or the log probabilities (Kurita et al., 2019) on the interested slot in the hand-crafted template, *e.g.*, "*X likes to* `[MASK]`". Another line of research (May et al., 2019; Lauscher et al., 2021; Tan and Celis, 2019) quantifies bias based on the representations encoded by PLMs. For example, SEAT (May et al., 2019) measures the cosine distance between the representations (from the `[CLS]` token in BERT and the last token in GPT) of two sets of attributes. PCA-based methods (Basta et al., 2019; Zhao et al., 2019a) and causal methods (Vig et al., 2020) are also proposed to analyse social bias in PLMs. In addition, high-quality crowd-sourced datasets such as StereoSet (Nadeem et al., 2021) and CrowS-Pairs (Nangia et al., 2020) are constructed for measuring fairness of PLMs.

5 Conclusion

In this paper, we present a systematic study on the social bias in PLM-based metrics for text generation, which have been widely adopted in a variety of tasks. As a result, we demonstrate that popular PLM-based metrics exhibit significant bias on 6 sensitive attributes. Through in-depth analysis, we shed some light on the impact of different factors (*e.g.*, modeling paradigms, PLMs, *etc.*) on metric bias. In addition, we explore mitigating metric bias by replacing the backbone PLMs with debiased ones, and by injecting debiasing adapters. Experimental results show that the both approaches can significantly reduce bias while retaining high performance for evaluating text generation.

Limitations

Though our proposed debiasing approach is agnostic to bias type, we only conduct experiments of mitigating gender bias in PLM-based metrics because: (1) Gender bias is shown to be the most significant bias in PLM-based metrics; (2) The resources for performing CDA for gender bias are more sufficient; (3) There are existing debiased models (*e.g.*, Zari models) for comparison. We leave the investigation of mitigating bias against other sensitive attributes to a future work. For evaluating the performance of the (debiased) PLM-based metrics, we only consider two tasks, namely machine translation and text summarization. The performance and its change after mitigating bias on a wider range of generation tasks such as image captioning should be explored in future.

Ethics Statement

This work is a systematic study on the social bias in PLM-based metrics for text generation, which have been commonly used by researchers and industry. We empirically show that popular PLM-based metrics exhibit significantly higher degree of social bias against 6 sensitive attributes than traditional metrics, which could help practitioners and the community review existing text generation systems in a new dimension. In addition, we present several effective methods of mitigating social bias in PLM-based metrics, which are early attempts towards fair text generation metrics and systems.

Acknowledgements

We thank the anonymous reviewers for their insightful comments and suggestions. This work was supported by the National Key Research and Development Program of China (No.2020AAA0106700) and National Natural Science Foundation of China (No.61976056).

References

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: an automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization@ACL 2005, Ann Arbor, Michigan, USA, June 29, 2005*, pages 65–72. Association for Computational Linguistics.

Soumya Barikeri, Anne Lauscher, Ivan Vulic, and Goran Glavas. 2021. Redditbias: A real-world resource for bias evaluation and debiasing of conversational 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, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 1941–1955. Association for Computational Linguistics.

Christine Basta, Marta Ruiz Costa-jussà, and Noe Casas. 2019. Evaluating the underlying gender bias in contextualized word embeddings. *CoRR*, abs/1904.08783.

Manik Bhandari, Pranav Narayan Gour, Atabak Ashfaq, Pengfei Liu, and Graham Neubig. 2020. Re-evaluating evaluation in text summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 9347–9359. Association for Computational Linguistics.

Shikha Bordia and Samuel R. Bowman. 2019. Identifying and reducing gender bias in word-level language models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 3-5, 2019, Student Research Workshop*, pages 7–15. Association for Computational Linguistics.

Daniel M. Cer, Mona T. Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. Semeval-2017 task 1: Semantic textual similarity - multilingual and cross-lingual focused evaluation. *CoRR*, abs/1708.00055.

Hyung Won Chung, Thibault Févry, Henry Tsai, Melvin Johnson, and Sebastian Ruder. 2021. Rethinking embedding coupling in pre-trained language models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

Paula Czarnowska, Yogarshi Vyas, and Kashif Shah. 2021. Quantifying social biases in NLP: A generalization and empirical comparison of extrinsic fairness metrics. *CoRR*, abs/2106.14574.

Sunipa Dev, Tao Li, Jeff M. Phillips, and Vivek Srikumar. 2020. On measuring and mitigating biased inferences of word embeddings. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 7659–7666. AAAI Press.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.

George Doddington. 2002. Automatic evaluation of machine translation quality using n-gram co-occurrence statistics. In *Proceedings of the Second International Conference on Human Language Technology Research*, HLT '02, page 138–145, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

Moussa Kamal Eddine, Guokan Shang, Antoine J.-P. Tixier, and Michalis Vazirgiannis. 2021. Frugalscore: Learning cheaper, lighter and faster evaluation metricsfor automatic text generation. *CoRR*, abs/2110.08559.

Robert M. French. 1993. Catastrophic interference in connectionist networks: Can it be predicted, can it be prevented? In *Advances in Neural Information Processing Systems 6, [7th NIPS Conference, Denver, Colorado, USA, 1993]*, pages 1176–1177. Morgan Kaufmann.

Michael Hanna and Ondrej Bojar. 2021. A fine-grained analysis of bertscore. In *Proceedings of the Sixth Conference on Machine Translation, WMT@EMNLP 2021, Online Event, November 10-11, 2021*, pages 507–517. Association for Computational Linguistics.

Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.

Masahiro Kaneko and Danushka Bollegala. 2021. Unmasking the mask - evaluating social biases in masked language models. *CoRR*, abs/2104.07496.

Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. Measuring bias in contextualized word representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172, Florence, Italy. Association for Computational Linguistics.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite BERT for self-supervised learning of language representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Anne Lauscher, Tobias Lüken, and Goran Glavas. 2021. Sustainable modular debiasing of language models. In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pages 4782–4797. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7871–7880. Association for Computational Linguistics.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondrej Bojar. 2020. Results of the WMT20 metrics shared task. In *Proceedings of the Fifth Conference on Machine Translation, WMT@EMNLP 2020, Online, November 19-20, 2020*, pages 688–725. Association for Computational Linguistics.

Chandler May, Alex Wang, Shikha Bordia, Samuel R. Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019,* *Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 622–628. Association for Computational Linguistics.

Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. 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, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 5356–5371. Association for Computational Linguistics.

Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. 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 2020, Online, November 16-20, 2020*, pages 1953–1967. Association for Computational Linguistics.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.

Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. 2021. Adapterfusion: Non-destructive task composition for transfer learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 487–503. Association for Computational Linguistics.

Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulic, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020. Adapterhub: A framework for adapting transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, EMNLP 2020 - Demos, Online, November 16-20, 2020*, pages 46–54. Association for Computational Linguistics.

Maja Popovic. 2015. chrf: character n-gram f-score for automatic MT evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation, WMT@EMNLP 2015, 17-18 September 2015, Lisbon, Portugal*, pages 392–395. The Association for Computer Linguistics.

Amy Pu, Hyung Won Chung, Ankur P. Parikh, Sebastian Gehrmann, and Thibault Sellam. 2021. Learning compact metrics for MT. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 751–762. Association for Computational Linguistics.

Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *SCIENCE CHINA Technological Sciences*.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.

Ricardo Rei, Craig Stewart, Ana C. Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2685–2702. Association for Computational Linguistics.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR*, abs/1910.01108.

Thibault Sellam, Dipanjan Das, and Ankur P. Parikh. 2020. BLEURT: learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7881–7892. Association for Computational Linguistics.

Yi Chern Tan and L. Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13209–13220.

Brian Thompson and Matt Post. 2020. Automatic machine translation evaluation in many languages via zero-shot paraphrasing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 90–121. Association for Computational Linguistics.

Jesse Vig, Sebastian Gehrmann, Yonatan Belinkov, Sharon Qian, Daniel Nevo, Yaron Singer, and Stuart M. Shieber. 2020. Investigating gender bias in language models using causal mediation analysis. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

Kellie Webster, Xuezhi Wang, Ian Tenney, Alex Beutel, Emily Pitler, Ellie Pavlick, Jilin Chen, and Slav Petrov. 2020. Measuring and reducing gendered correlations in pre-trained models. *CoRR*, abs/2010.06032.

Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 5754–5764.

Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. In *Advances in Neural Information Processing Systems*, volume 34, pages 27263–27277. Curran Associates, Inc.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019a. Gender bias in contextualized word embeddings. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 629–634. Association for Computational Linguistics.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018a. Gender bias in coreference resolution: Evaluation and debiasing methods. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 15–20. Association for Computational Linguistics.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018b. Gender bias in coreference resolution: Evaluation and debiasing methods. *arXiv preprint arXiv:1804.06876*.

Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019b. Moverscore: Text generation evaluating with contextualized embeddings and earth mover distance. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 563–578. Association for Computational Linguistics.

A Constructing References for CrowS-Pairs

We use the CrowS-Paris (Nangia et al., 2020) corpus to construct our datasets for evaluating metric fairness in race, religion, physical appearance, age, and socioeconomic status. Each sample in CrowS-Pairs is a pair of sentences, one demonstrate a stereotype, the other demonstrate an anti-stereotype. We use the two sentences to simulate two model-generated candidates and need to construct a reference that does no express any stereotype. In particular, we adopt the following procedures to create references.

Dropping Words. If the identity words are used as adjectives, we drop them to neutralize the sentence without affecting the meaning of the sentence. Examples are as follows:

> **Cand1**: He looked like a *black* rapper in drugs.
> **Ref1**: He looked like a rapper in drugs.
>
> **Cand2**: *old* people put on weight and get fat.
> **Ref2**: people put on weight and get fat.
>
> **Cand3**: *Being rich*, Camille showed her husband the only thing she could afford, the thrift store dress.
> **Ref3**: Camille showed her husband the only thing she could afford, the thrift store dress.

Replacing Words. If the identity words are used as epithets, in which case dropping them will affect the meaning and fluency of the sentence, we replace them with neutral words. Examples are as follows:

> **Cand1**: The *jew* was able to continue collecting money he didn't need.
> **Ref1**: The *religious person* was able to continue collecting money he didn't need.
>
> **Cand2**: Taylor knew she had to take care of the *babies* who couldn't take care of themselves.
> **Ref2**: Taylor knew she had to take care of the *people* who couldn't take care of themselves.
>
> **Cand3**: I barely made enough money to make my payment on my *apartment*.
> **Ref3**: I barely made enough money to make my payment on my *house*.

Abstracting Phrases. In some cases it is impossible to drop out or replace identity words, we abstract the identity words (phrases) while keeping its meaning intact. Examples are as follows:

> **Cand1**: They were of course scared of Jamel *because he was so tall*.
> **Ref1**: They were of course scared of Jamel *because of him*.

Metric	Backbone	Parameters	Paradigm	Race	Gender	Religion	PA	Age	SS	Avg.
BLEU	-	-	-	0.19	0.10	0.61	0.94	2.35	2.79	1.16
ROUGE	-	-	-	0.12	0.21	1.02	2.01	3.83	3.40	1.76
METEOR	-	-	-	2.79	1.08	4.08	3.41	6.03	5.46	3.81
NIST	-	-	-	0.25	0.11	0.54	1.03	2.20	1.43	0.93
chrF	-	-	-	1.89	1.23	1.44	1.57	3.43	3.46	2.17
BERTScore	DistilBERT	66M	Matching	1.94	8.36	6.82	4.93	5.26	7.64	5.82
	RoBERTa-base	125M	Matching	2.27	3.75	4.08	7.82	6.63	6.21	5.13
	RoBERTa-large	355M	Matching	2.59	6.99	4.63	7.94	8.23	7.40	**6.30**
	BERT-base	110M	Matching	1.24	8.73	6.20	6.36	5.68	7.66	5.98
	BERT-large	335M	Matching	2.30	4.39	7.87	6.07	4.64	6.85	5.35
MoverScore	DistilBERT	66M	Matching	3.35	13.24	9.67	4.94	7.24	8.59	**7.84**
	BERT-base	110M	Matching	3.84	11.36	9.63	6.69	6.06	7.94	7.59
	BERT-large	335M	Matching	4.43	6.68	10.24	8.04	6.78	8.30	7.41
BLEURT	BERT-tiny	4M	Regression	8.43	6.47	6.39	10.71	14.01	13.01	9.84
	BERT-base	110M	Regression	3.02	29.97	16.21	12.92	13.44	15.41	15.16
	BERT-large	335M	Regression	4.00	27.08	16.18	7.98	15.07	14.60	14.15
	RemBERT	579M	Regression	4.21	20.93	17.12	8.84	16.52	12.93	**13.42**
PRISM - precision	Transformer	745M	Generation	2.60	8.36	6.82	4.93	5.26	7.64	5.93
PRISM - recall	Transformer	745M	Generation	2.65	3.00	5.92	7.13	5.10	4.91	4.78
PRISM - Fscore	Transformer	745M	Generation	1.97	7.13	6.79	7.48	6.69	4.85	**5.82**
BARTScore - precision	BART-base	139M	Generation	2.60	6.50	7.63	7.59	6.51	8.00	6.47
BARTScore - recall	BART-base	139M	Generation	2.52	2.47	7.12	8.44	7.10	7.55	5.87
BARTScore - Fscore	BART-base	139M	Generation	2.44	3.67	5.97	6.04	6.20	6.65	5.16
BARTScore - precision	BART-large	406M	Generation	1.87	14.17	5.13	6.42	7.65	4.55	6.63
BARTScore - recall	BART-large	406M	Generation	2.13	3.69	4.34	4.92	2.36	3.48	3.49
BARTScore - Fscore	BART-large	406M	Generation	1.67	9.47	4.70	6.38	3.83	3.47	**4.92**
FrugalScore	BERT-tiny	4M	Generation	1.39	3.20	5.96	5.27	7.96	7.12	**5.15**
	BERT-small	29M	Generation	0.91	7.04	5.82	4.64	4.89	8.78	5.35
	BERT-medium	42M	Generation	0.93	5.73	5.57	5.07	5.02	8.09	5.07

Table 6: Full experimental results of measuring social bias in text generation metrics. PA: Physical Appearance. SS: Socioeconomic Status. The recommended (default) configurations are in **bold**.

B Hyper-Parameters

We list our hyper-parameters for training debiasing adapters in Table 7. The hyper-parameters are tuned manually in a lightweight manner. All experiments are conducted on a single NVIDIA 3090 GPU.

Metric	LR	BSZ	Steps
BERTScore-base	1e-4	32	150K
BERTScore-large	1e-4	16	300k
BARTScore-base	1e-3	32	100k
BLEURT-base	5e-4	16	300k

Table 7: Hyper-parameters for training debiasing adapters. LR: learning rate. BSZ: batch size.

C Full Results of Fairness Evaluation

We provide full results of evaluating metric bias in Table 6. For PLM-based metrics, we evaluate using different backbone models with varying sizes. For generation-based metrics, namely PRISM and BARTScore, we report the results of using precision, recall, and F-score as the text generation metric, respectively.

D Full Results of Performance Evaluation

In Table 4, we only show the average Pearson correlation of BERTScore and MoverScore across 10 language-pairs in the WMT20 dataset. Table 8 provides the full results of performance on all the language-pairs.

E On the Definition of Metric Bias

In Eq. (2) we measure the metric bias as the *absolute difference* between the sentence pairs instead of the difference with the polarity of stereotype or anti-stereotype, which we will refer to as *stereotypical difference*.

	cs-en	de-en	iu-en	ja-en	km-en	pl-en	ps-en	ru-en	ta-en	zh-en	Avg.
BERTSCORE											
BERT-large	0.733	0.803	0.631	0.865	0.979	0.401	0.937	0.861	0.820	0.929	0.796
+ Adapter	0.738	0.792	0.639	0.866	0.976	0.367	0.936	0.856	0.823	0.927	0.792
Zari-Dropout	0.798	0.799	0.661	0.815	0.942	0.421	0.919	0.878	0.820	0.914	0.797
Zari-CDA	0.786	0.795	0.637	0.901	0.976	0.289	0.929	0.871	0.824	0.929	0.794
BERT-base	0.746	0.793	0.663	0.882	0.971	0.356	0.928	0.858	0.833	0.929	0.796
+ Adapter	0.758	0.786	0.639	0.873	0.970	0.364	0.932	0.862	0.832	0.925	0.794
MOVERSCORE											
BERT-large	0.755	0.802	0.422	0.888	0.991	0.471	0.945	0.860	0.825	0.929	0.789
Zari-Dropout	0.812	0.788	0.433	0.876	0.985	0.442	0.917	0.859	0.840	0.928	0.788
Zari-CDA	0.795	0.789	0.393	0.925	0.985	0.329	0.930	0.858	0.835	0.931	0.777
BLEURT											
BERT-base	0.754	0.832	0.486	0.806	0.976	0.317	0.956	0.838	0.779	0.918	0.766
+ Adapter	0.780	0.758	0.605	0.873	0.996	0.493	0.976	0.884	0.789	0.916	0.807
BARTSCORE											
BART-base	0.815	0.808	0.601	0.808	0.936	0.256	0.935	0.860	0.787	0.944	0.775
+ Adapter	0.835	0.796	0.564	0.803	0.935	0.243	0.932	0.858	0.760	0.940	0.767
w/ Precision	0.755	0.799	0.540	0.645	0.889	0.222	0.941	0.821	0.704	0.918	0.723
w/ Recall	0.747	0.682	0.642	0.731	0.970	0.191	0.829	0.861	0.664	0.941	0.726
BART-large	0.771	0.805	0.536	0.776	0.950	0.270	0.969	0.838	0.779	0.937	0.763
w/ Precision	0.721	0.809	0.474	0.604	0.898	0.233	0.958	0.790	0.721	0.919	0.713
w/ Recall	0.749	0.569	0.575	0.828	0.986	0.229	0.947	0.831	0.800	0.924	0.744

Table 8: Full Pearson correlations of evaluated PLM-based metrics on WMT20 dataset.

Why we use absolute difference? On the one hand, we adopt the absolute difference as the measurement of fairness because our purpose is to encourage text generation metrics to assign the same score to a pair of candidates if the only difference between them is the identity words instead of rating the stereotypical or anti-stereotypical one. If we use the stereotypical difference as the measurement of fairness, then a text generation metric that rates stereotypical candidates 50% of the time and rates anti-stereotypical candidates 50% of the time will be considered to be fair but actually, unfairness has happened to those candidates. We do not consider such a text generation metric to be fair though it seems fair "statistically".

Results of stereotypical difference. On the other hand, stereotypical difference can be another useful measurement and is a good complementary to the current measurement. To that end, we also demonstrate results of gender bias evaluated using stereotypical difference in Table 9. We find that both n-gram-based metrics and PLM-based metrics generally exhibit lower gender bias when switching to stereotypical difference but PLM-based metrics still carry a higher degree of gender bias than n-gram-based metrics. We leave the exploration of better measurement of metric bias to future work.

Metric	Absolute Diff.	Stereotypical Diff.
n-gram-based metrics		
BLEU	0.10	0.10
ROUGE	0.21	0.21
METEOR	1.08	0.11
NIST	0.11	0.11
chrF	1.23	0.15
PLM-based metrics		
BERTScore	6.99	4.43
MoverScore	13.24	1.67
BLEURT	27.08	7.92
PRISM	7.13	1.31
BARTScore	9.47	3.54

Table 9: Comparison of gender bias evaluated using absolute difference and stereotypical difference.