

Article

What Is Your Favorite Gender, MLM? Gender Bias Evaluation in Multilingual Masked Language Models

Jeongrok Yu *, Seong Ug Kim, Jacob Choi  and Jinho D. Choi * 

Department of Computer Science, Emory University, Atlanta, GA 30322, USA; jacobjch@usc.edu (J.C.)

* Correspondence: jyu673@gatech.edu (J.Y.); jinho.choi@emory.edu (J.D.C.)

Abstract: Bias is a disproportionate prejudice in favor of one side against another. Due to the success of transformer-based masked language models (MLMs) and their impact on many NLP tasks, a systematic evaluation of bias in these models is now needed more than ever. While many studies have evaluated gender bias in English MLMs, only a few have explored gender bias in other languages. This paper proposes a multilingual approach to estimating gender bias in MLMs from five languages: Chinese, English, German, Portuguese, and Spanish. Unlike previous work, our approach does not depend on parallel corpora coupled with English to detect gender bias in other languages using multilingual lexicons. Moreover, a novel model-based method is presented to generate sentence pairs for a more robust analysis of gender bias. For each language, lexicon-based and model-based methods are applied to create two datasets, which are used to evaluate gender bias in an MLM specifically trained for that language using one existing and three new scoring metrics. Our results show that the previous approach is data-sensitive and unstable, suggesting that gender bias should be assessed on a large dataset using multiple evaluation metrics for best practice.

Keywords: bias evaluation; multilingual bias benchmark; multilingual language models



Citation: Yu, J.; Kim, S.U.; Choi, J.; Choi, J.D. What Is Your Favorite Gender, MLM? Gender Bias Evaluation in Multilingual Masked Language Models. *Information* **2024**, *15*, 549. <https://doi.org/10.3390/info15090549>

Academic Editor: William D. O'Grady

Received: 6 August 2024

Revised: 2 September 2024

Accepted: 3 September 2024

Published: 7 September 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The advent of transformer models [1] and subsequent development of contextualized embedding encoders [2,3] have led to the widespread deployment of large language models for crucial societal tasks [4]. However, the use of such models has brought to light critical concerns regarding bias [5–7]. Despite efforts to enhance masked language models (MLMs) [8,9], which are adapted to pre-train the transformers for language understanding by predicting masked tokens from their context, the use of sophisticated models and extensive datasets has intensified worries about bias in MLMs.

The ubiquity of language models in society has sparked a growing interest in detecting and mitigating their inherent biases. Detecting gender disparities in language technologies has gained traction across multiple domains, as evidenced by studies on identifying human-like biases in a transformer encoder [10] and creating benchmarks for testing gender bias [11]. These works have recognized undesirable biases in language models, leading to more studies on revealing bias in embeddings [12]. Indeed, language models can have bias issues that must be addressed to ensure equitable and inclusive outcomes [13,14].

There has been pioneering work on evaluating gender bias in static word embeddings using word analogies [15], while the evaluation of gender bias in contextualized word embeddings has largely focused on monolingual models. Zhao et al. [16] analyzed word pairs to differentiate which words contain gender information/bias. Liang et al. [17] examined the MLMs' tendency to predict gendered pronouns associated with certain occupations. There have also been innovative debiasing methods for MLMs, such as re-balancing the corpus by switching bias attribute words [18] or using derivation to normalize sentence vectors from MLMs [19]. Nonetheless, there remains a gap in research on gender bias in multilingual MLMs, which has not been explored to the same extent.

This presents a crucial area for research to ensure that multilingual language models are free from harmful biases.

This paper improves the gender bias evaluation in multilingual MLMs by addressing the limitations of previous work. Section 3 identifies such limitations and presents an enhanced method to generate sentence pairs for gender bias evaluation in MLMs. Section 4 provides multilingual lexicons to detect sentences with gendered words from five languages. Section 5 compares the performance of our method to that of existing methods. Finally, Section 6 shows that our method retains more data from the target corpus and is more consistent than the other methods in evaluating gender bias, especially when the data are skewed in the gender distribution. Our main contributions are as follows:

- We create multilingual gender lexicons to detect sentences with gendered words in Chinese, English, German, Portuguese, and Spanish without relying on parallel datasets, which enables us to extract more diverse sets of gendered sentences and facilitate more robust evaluations.
- We present two novel metrics and methods that provide a rigorous approach to comparing multilingual MLMs and datasets. Our approach ensures meaningful and fair comparisons, leading to more reliable and comprehensive assessments of gender bias in multilingual MLMs.

2. Related Work

The issue of bias in large language models has garnered significant attention in recent years, prompting several studies that aim to assess and address the presence of bias in MLMs. Nangia et al. [20] introduced CrowS-Pairs, a dataset consisting of single sentences with masked attribute words, aiming to assess potential social bias in terms of race, gender, and religion in MLMs. Nadeem et al. [21] adopted a similar approach by masking modified tokens to measure bias. These studies were confined to English, however, and lacked a precise definition of the term “bias”, resulting in ambiguity and assumptions about its meaning [22].

Ahn and Oh [23] introduced a novel approach to evaluating bias in MLMs using pairs of sentences with varying degrees of masking. They proposed a new metric, the Categorical Bias score, showing the variance of log-likelihood interpreted as an effect size of the attribute word. In addition, they analyzed ethnic bias across six languages, making an effort to generalize bias evaluation. However, their method still required human-written sentences with bias annotations, which is a limitation in terms of capturing the natural usage of language and can be exploited when the model finds a simple loophole around the set of rules [24]. Moreover, there has been some research that has utilized evaluation metrics to look into debiasing methods. Our work requires minimal annotation and is evaluated on a real dataset rather than a contrived one.

A few studies have attempted to evaluate bias in multilingual MLMs. Kaneko and Bollegala [25] assessed bias by utilizing a set of English words associated with males and females and computing their likelihoods. Kaneko et al. [26] used a parallel corpus in English and eight other languages, where bias was annotated solely in English, to evaluate gender bias in the target language models. Our approach is distinguished in that it does not rely on a parallel corpus, making it language-independent.

3. Methodology

3.1. Multilingual Bias Evaluation

Kaneko et al. [26] proposed a Multilingual Bias Evaluation (MBE) score to assess the gender bias in MLMs across many languages using a multilingual corpus with English translations. To ensure a fair comparison with previous work, we adopt the MBE score as our baseline metric. The MBE score employs a three-step approach to detect potential gender bias.

First, MBE extracts English sentences containing one male or female noun within a sentence. The male and female nouns used to pick out sentences are from a set of gendered

nouns presented by Bolukbasi et al. [15] and common first names from Nangia et al. [20]. MBE selects the corresponding sentence in the target language, under the premise that the parallel sentence also contains gendered terms. The extracted sentences with respective single gender words are categorized into \mathcal{T}_f and \mathcal{T}_m , representing sets of female and male sentences.

For each sentence, $S = [w_1, \dots, w_n] \in (\mathcal{T}_f \cup \mathcal{T}_m)$, where w_i is the i th token in S , MBE estimates its All Unmasked Likelihood with Attention (AULA; Kaneko and Bollegala [25]). Given all tokens in S , except for w_i , and pre-trained parameters θ , the likelihood of an MLM predicting w_i is measured by $P_{\text{MLM}}(w_i | S \setminus w_i; \theta)$. AULA is then computed by summing the log-likelihoods of all w_i multiplied by the averages of the multi-head attention weights α_i associated with w_i , as in Equation (1):

$$A(S) = \frac{1}{|S|} \sum_{\forall w_i \in S} \alpha_i \cdot \log P_{\text{MLM}}(w_i | S \setminus w_i; \theta) \quad (1)$$

Thus, AULA estimates the relative importance of individual words in the sentence by the MLM. Finally, the MBE score is measured by comparing all sentences between \mathcal{T}_f and \mathcal{T}_m , as in Equation (2):

$$\frac{\sum_{\forall S_m \in \mathcal{T}_m} \sum_{\forall S_f \in \mathcal{T}_f} \gamma(S_m, S_f) \cdot \mathbb{I}(S_m, S_f)}{\sum_{\forall S_m \in \mathcal{T}_m} \sum_{\forall S_f \in \mathcal{T}_f} \gamma(S_m, S_f)} \quad (2)$$

$\gamma(S_f, S_m)$ is the cosine similarity score between sentence embeddings of S_f and S_m . The indicator function $\mathbb{I}(S_m, S_f)$ returns 1 if $A(S_m) > A(S_f)$; otherwise, 0. Therefore, the MBE score shows the percentage of male sentences that the MLM generates over female sentences in the parallel corpus.

3.2. Strict Bias Metric

A significant drawback of MBE is its potential lack of rigor in comparisons. One key limitation lies in the use of AULA to calculate sentence likelihoods that can lead to tokens unrelated to gendered words exerting undue influence on the score. This is even more problematic when comparing AULAs for sentences that are notably dissimilar.

To address unexpected measurement errors in evaluating MLM bias, we propose the strict bias metric (SBM), which solely compares likelihoods between parallel sentences that differ only by gendered words, minimizing potential confounding factors. As the number of comparisons in AULAs decreases when using SBM, it allows for a more focused and targeted analysis so that SBM guarantees a meaningful assessment by capturing the likelihood differences incurred only by variations in gendered words. The SBM score is measured as in Equation (3):

$$\frac{\sum_{(S_m, S_f) \in (\mathcal{T}_m \times \mathcal{T}_f)} \mathbb{I}(S_m, S_f)}{|\mathcal{T}_m|} \quad (3)$$

Since SBM requires the identification of the specific gender words, we create a multilingual lexicon that provides sets of gendered words in five languages (Section 4.1). Moreover, since SBM needs sentence pairs that differ only by gender words, we present two methods, lexicon-based (Section 3.3) and model-based (Section 3.4), to generate such sentence pairs using solely monolingual datasets.

3.3. Lexicon-Based Sentence Generation

The lexicon-based sentence generation (LSG) first identifies every sentence containing a single gender word and constructs its counterpart by replacing the gender word with its opposite-gendered word, which is provided in the lexicon (Section 4.1). For the sentence,

“The *waitress* came over”, LSG generates the new sentence “The *waiter* came over” by replacing the gender word (Figure 1). SBM is then computed between these two sentences, taking into account the relative importance of those gender words.

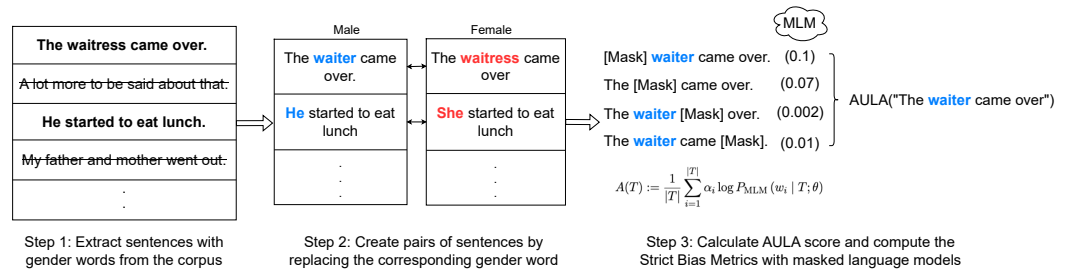


Figure 1. An overview of the lexicon-based sentence extraction (Section 3.3).

Sentences containing multiple gender words are excluded in this approach because it is ambiguous to identify them as either male- or female-gendered sentences. For example, a sentence “The *actor* fell in love with the *queen*” is considered neither a male nor female sentence because it contains words for both genders and masking one of them for an MLM would make this sentence the opposite gender. Including such sentences would undermine the reliability of the SBM score. Thus, sentences of this nature are excluded from our experiments to ensure the robustness of bias evaluation.

The corpus may have an imbalanced distribution of sentences with male and female words that can result in a bias in the overall data representation. To address this issue, we extract an equal number of male and female sentences, such that $|\mathcal{T}_m| = |\mathcal{T}_f|$. It is worth mentioning that we tried to discard sentences exhibiting strong contextual biases towards a specific gender from our evaluation set, as they may introduce inherent biases and compromise the results. We achieved this by manually going over the sentences to pick out verbs that might be only applicable to one gender and disposing of those sentences (e.g., The *princess* gave birth to twins).

3.4. Model-Based Sentence Generation

Figure 2 describes an overview of our model-based sentence generation (MSG). Similar to LSG, it begins by collecting sentences with only one gender word. It then masks the gender word in each sentence and employs an MLM to predict the most likely male and female words based on the context (e.g., the first row in Figure 2). This enables MSG to select the most probable gender words for every sentence by MLM’s predictions, ensuring that the comparisons are conducted on model-derived sentence pairs.

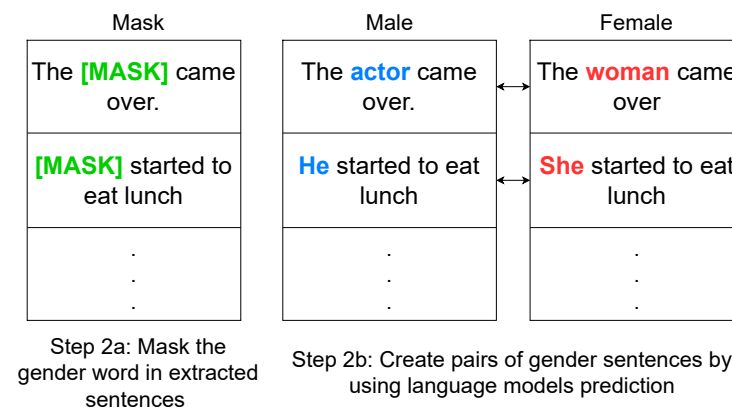


Figure 2. An overview of the model-based sentence extraction (Section 3.4) using the examples in Figure 1.

In theory, MSG can always generate both male and female words for any masked word by taking words from the corresponding gender sets whose MLM scores are the highest

among the others in the sets. However, in practice, MSG may fail to generate the sentence pair because the MLM does not predict “meaningful” gender words with high confidence. In this case, MSG adapts LSG to handle the missing sentence. For the second-row example in Figure 2, the MLM effectively predicts the female word ‘She’ but does not predict any male word with high confidence, in which case, it takes the opposite gendered word ‘He’ as its lexicon-based counterpart. On the other hand, sentences for which the MLM fails to predict both male or female words with high scores are discarded. Once all sentence pairs are created, SBM is applied to assess potential bias in the MLM towards any particular gender.

For our experiments, a threshold of 0.01 is used to determine whether MLM’s prediction confidence is high. This threshold is observed by manual inspection of MLM predictions in five languages. In addition, we use the threshold to configure the top- k reliable predictions for MLMs. To determine the optimal value for k , an iterative analysis is conducted with k ranging from 1 to 15, assessing how many predictions above the threshold are covered by the top k predictions. Our findings indicate that $k = 10$ is the optimal choice, as the convergence rate of the model exhibits only marginal changes beyond this cutoff point (Figure 3).




Figure 3. The proportions of sentences containing gender words whose confidence scores are higher than the threshold 0.01 covered by the $k = [1, 15]$ predictions.

3.5. Direct Comparison Bias Metric

SBM quantifies the extent to which MLMs prefer to predict male words over female words given the neutral context. Since MLMs provide confidence scores for the predicted words, it is also possible to make word-level comparisons. Thus, we propose the Direct Comparison Bias Metric (DBM), which compares the scores of the predicted gender words for every (predictable) sentence.

Given a sentence containing one gender word, an MLM is employed to predict male and female words. If the scores for both male and female word predictions are below the threshold (Section 3.4), the sentence is excluded as the comparison, in this case, is not meaningful. Then, the DBM is measured by replacing the indicator function $\mathbb{I}(S_m, S_f)$ in Equation (3) with $\mathbb{I}_d(w_m, w_f)$, where $w_m \in S_m$ is the predicted male word and $w_f \in S_f$ is the predicted female word. \mathbb{I}_d returns 1 if the MLM’s prediction score of w_m is higher than that of w_f ; otherwise, 0.

4. Data Preparation

4.1. Multilingual Gender Lexicon

We use the male and female word list provided by Bolukbasi et al. [15] to construct our English gender word set. This set consists of pairs of words, with each gender word

having its counterpart of the opposite gender, which makes it suitable for LSG (Section 3.3) and MSG (Section 3.4). Unlike Kaneko et al. [26], we exclude common first names from the CrowS-Pairs dataset [20] when making our lexicon. The aim of this is to avoid potential issues related to transliteration during name translation, which can lead to MLMs treating the first names as phonetically similar words in the target language.

Our Multilingual Gender Lexicon (MGL) is compiled by translating gender words in the English set using automatic systems, such as Bing (Microsoft Bing Translator: <https://www.bing.com/translator> (accessed on 6 August 2024)), DeepL (DeepL Translate: <https://www.deepl.com/translator> (accessed on 6 August 2024)), and Google (Google Translate: <https://translate.google.com> (accessed on 6 August 2024)), into the following eight languages: Arabic, Chinese, German, Indonesian, Japanese, Portuguese, Russian, and Spanish. If the majority of the reviewers find a translation to be unnatural or gender-neutral, both the translation and its counterpart are excluded from MGL. For example, the pronoun 'Sie' in German has two meanings: 'she' and 'you' (honorific). Although 'sie' commonly refers to a female and capitalized 'Sie' refers to 'You', it was removed from our lexicon to leave no room for ambiguity regarding its connotation of gender information. This meticulous process ensures that MGL contains accurate and natural translations of only gendered words in each language.

4.2. MGL Validation

The coverage of MGL is assessed by comparing the words in parallel sentences extracted from the TED corpus (2020 v1; Reimers and Gurevych [27]) (TED2020 v1: <https://opus.nlpl.eu/TED2020.php> (accessed on 6 August 2024)). The TED corpus consists of about 4000 TED talks that comprise a total of 427,436 English sentences. Many of the English sentences have been translated into 100+ languages by certified translators.

For the validation, a set \mathcal{E}_r of 11,000 English sentences is randomly sampled from the TED corpus. The sentences in \mathcal{E}_r are checked against the English gender word set (Section 4.1) to create another set \mathcal{E}_g of English gendered sentences. Next, for every target language ℓ , a set \mathcal{G}^ℓ is created by finding ℓ 's translations of the sentences in \mathcal{E}_g from the corpus. Not all sentences in \mathcal{E}_g may come with translations, such that $|\mathcal{G}^\ell| \leq |\mathcal{E}_g|$. Finally, a new set \mathcal{G}_g^ℓ of gendered sentences is created where each sentence in \mathcal{G}_g^ℓ includes the translated gender word according to the target-language gender word (Section 4.1). Table 1 illustrates the results from this validation, including the coverage percentages of MGL for the eight target languages.

MGL's coverage rates in gender words are below 0.5 for the following four languages: Indonesian, Russian, Japanese and Arabic. Many gender words in English are translated into gender-neutral words in Indonesian because it would sound unnatural to use gender-specific translations [28]. Both Russian and Arabic are morphologically rich languages with extensive inflectional/derivational systems [29,30], which contribute to their low coverage rates. Japanese is a pro-drop language that allows the omission of the subject in a sentence, which makes identifying gender words challenging because the subject pronouns, such as 'he' and 'she', are often dropped in natural discourse.

It is important to note that the assumption made by Kaneko et al. [26] that the gender information in an English sentence containing a gender word is retained in parallel sentences in the target languages, even if the corresponding gender words do not exist in the parallel sentences, may not hold true for certain languages mentioned above (e.g., Indonesian, Japanese). Thus, our bias evaluation uses only languages with high coverage rates: English, German, Spanish, Portuguese, and Chinese.

Table 1. The coverage rates of sentences containing gender words from MGL across the eight target languages.

Language	$ \mathcal{G}^\ell $	$ \mathcal{G}_g^\ell $	Coverage (%)
German	1226	1124	91.7
Spanish	1380	1125	81.5
Portuguese	1206	928	76.9
Chinese	1325	997	75.2
Indonesian	671	312	46.5
Russian	1289	583	45.2
Japanese	1288	466	36.6
Arabic	1327	252	19.0

4.3. Sentence Pair Generation

Two sets of sentence pairs are created to evaluate the gender bias of MLMs, one using LSG and the other using MSG, in English as well as the top four languages in Table 1, all of which contain sufficient numbers of gendered words available in MGL. For our experiments, BERT-based language-specific transformer encoders are used for English [2], German [31], Spanish [32], Portuguese [33], and Chinese [34]. The statistics of these datasets in comparison to the one used by Kaneko et al. [26] are presented in Table 2.

Table 2. The statistics of gender bias evaluation datasets.

Language	Kaneko _{org}	Kaneko _{all}	LSG	MSG	Total
English	—	39,040	25,993	28,112	34,970
Chinese	6800	36,270	22,196	22,616	30,547
German	4700	26,639	32,436	29,667	33,154
Portuguese	5700	29,975	24,608	31,670	36,072
Spanish	7100	37,808	76,972	96,995	114,168

Kaneko_{org} refers to the number of gendered sentences used by Kaneko et al. [26]. **Kaneko_{all}** is the number of all gendered sentences extracted from the TED corpus utilizing the same method as that described in Kaneko et al. [26]. Note that the previous work used only a subset of **Kaneko_{all}** for their evaluations. **LSG & MSG** shows the number of sentence pairs generated by lexicon-based (Section 3.3) and model-based (Section 3.4) methods, respectively; the numbers in these columns should be doubled for a fair comparison with the numbers in **Kaneko_{*}**. **Total** indicates the total number of sentences extracted from the TED corpus by using MGL (before discarding any sentences for balancing).

All sentences consisting of single, animate, gendered referents are extracted from the TED corpus by using MGL. For LSG, an equal number of male and female sentences are extracted for the creation of each dataset in Table 2 to minimize potential contextual bias that arises from imbalanced distributions between male and female sentences (Section 3.3). However, this constrains the extracted sentences to the size of the smaller gender category. For example, if there exist 100 male sentences but only 50 female sentences, half of the male sentences are discarded to match the number of female sentences. Consequently, the total number of extracted sentences becomes 100, not 150. In this case, even if there is contextual bias present in those sentences, the balanced number of sentences ensures that it does not significantly affect the overall evaluation results.

For MSG, on the other hand, the MLM generates both male and female sentences for the majority of extracted sentences. Although MLM sometimes generates one-sided predictions, the ratio between male and female sentences produced by MSG is often more balanced than that of LSG. For the previous example, where 150 sentences are extracted (100 male and 50 female), assume that the MLM generates both male and female sentences for 60, only male sentences for 40, only female sentences for 30, and no sentences for 20 instances. MSG discards the 20 cases for which the MLM generated no sentences and

10 (out of the 40) male-only sentences, so its count matches the one for the female-only sentences, resulting in a retention of 120 sentences, where the ratio between male and female sentences is even. Table 3 shows the proportions among the extracted sentences MSG generated for both genders, one gender, and no gender.

Table 3. The proportions of sentences extracted using MGL, for which MSG generates sentences for both genders, one gender, and none.

Language	Both	One	None
English	63.8%	16.6%	19.6%
Chinese	59.1%	14.9%	26.0%
Spanish	51.9%	33.1%	15.0%
Portuguese	43.0%	44.8%	12.2%
German	30.7%	58.8%	10.5%

For languages such as German, Portuguese, and Spanish, it is important to ensure grammatical correctness in the sentences generated by LSG and MSG. This is because these languages have gendered articles, adjectives, demonstratives, possessives, and attributive pronouns that need to agree with the genders of the generated words. For German, 19.85% of the extracted sentences contain such gendered components. To address this, we add mappings between gender words and their dependent gendered components to MGL for those languages. We then use the mappings to replace gendered components in the generated sentences accordingly. For example, in a German sentence “ein guter Mann” (a good man), when Mann (*man*) is replaced with Frau (*woman*), the article *ein* is also replaced with *eine* and the adjective *guter* is replaced with *gute* using the mapping provided in MGL. This ensures that the generated sentences maintain agreement between gender words and their associated components.

5. Experiments

We utilize the TED parallel corpus (Section 4.2) to build a comprehensive corpus for the evaluation of gender bias in the transformer-based multilingual language models (Section 4.3). For evaluation, the transformer implementation of Wolf et al. [35] is employed. All experiments are conducted on an Apple M1-Pro Chip with a 14-core GPU. The entire evaluation process, which includes processing all sentence pairs in every language, is completed in about two hours with efficient performance.

5.1. Multilingual Bias Evaluation on Kaneko*

Upon replicating Kaneko et al. [26], we observe that the previous study utilized only $\approx 25\%$ of the sentences in the TED corpus (Table 2). For a more reliable evaluation, we reconstruct these datasets by following their methods using all sentences in the corpus and compute the MBE scores (Section 3.1). Table 4 shows the results from Kaneko et al. [26] (Kaneko_{org}), as well as the results evaluated on our reconstructed datasets (Kaneko_{all}) (The evaluation on Kaneko_{all} is conducted five-fold, similar to LSG, for the same reason as explained in Section 5.2).

Surprisingly, while the English MLM shows bias towards male terms in Kaneko_{all}, MLMs in all the other four languages exhibit bias towards female terms, which contradicts the findings in Kaneko_{org}. The discrepancy between the results obtained from Kaneko_{org} and Kaneko_{all} highlights the limitations of working with smaller datasets.

Table 4. Left Group: Multilingual Bias Evaluation (MBE) scores evaluated on the sub-sampled sentences (Kaneko_{org}) and all sentences in the TED corpus (Kaneko_{all} ; \pm : standard deviation). Middle Group: The strict bias metric (SBM) scores achieved by LSG and MSG, as well as the Direct Comparison Bias Metrics (DBM) scores obtained by MSG. Right Group: The distributions of (and the ratios between) male and female sentences in the TED corpus (in %).

Language	Kaneko_{org}	Kaneko_{all}	LSG	MSG	DBM	Male/Female (Ratio)
English	—	52.07 (± 1.34)	50.39 (± 0.28)	45.49	75.18	62.83/37.17 (1.69:1)
Chinese	52.86	46.67 (± 0.55)	46.42 (± 0.68)	53.15	89.62	63.67/36.33 (1.75:1)
German	54.69	45.78 (± 1.72)	52.31 (± 0.64)	55.43	44.72	48.92/51.08 (0.96:1)
Portuguese	53.07	46.70 (± 0.81)	51.77 (± 0.44)	61.04	73.36	65.89/34.11 (1.93:1)
Spanish	51.44	48.52 (± 1.04)	41.68 (± 0.76)	50.74	72.34	66.29/33.71 (1.97:1)

5.2. Strict Bias Metric for LSG and MSG

To assess gender bias using LSG, for each language, we create five folds of evaluation datasets by randomly truncating the larger gender set (Section 4.3), while keeping the sentences in the smaller gender set unchanged across all folds. Hence, sentences in the larger gender set vary across the different folds. Table 4 presents the SBM scores (Section 3.2) of this evaluation for the five languages. Our results reveal that English, German, and Portuguese MLMs are biased towards males, with the scores greater than 0.5. In contrast, Chinese and Spanish MLMs depict a bias towards females. These trends are consistent across all five evaluation folds; thus reinforcing the reliability of our findings.

For MSG, the SBM scores are also used to assess gender bias. Unlike LSG, only a small portion of the larger gender sets get truncated using this method; thus, we create a single fold for the MSG evaluation. Our results show that the English MLM has a bias towards females, although the other MLMs show a bias towards males. Interestingly, the findings from LSG and MSG disagree in terms of the bias directions for English, Chinese, and Spanish. Several factors may contribute to this, such as differences in the pre-trained MLMs or characteristics of the languages, which we will explore in the future.

5.3. Direct Comparison Bias Metric for MSG

Finally, we assess gender bias using the DBM metric (Section 3.5) on the MSG datasets. The scores obtained using this metric are considerably more extreme than those from MBE and SBM. Notice that the DBM scores align with the gender distributions shown in the TED corpus (Table 4), implying that using DBM to evaluate gender bias of MLMs on a corpus with a significant gender imbalance can lead to unreliable results, although they are suitable for quantifying the bias in the evaluation corpus itself.

6. Analysis

The discrepancy in results between Kaneko_{org} and Kaneko_{all} , as discussed in Section 4.1, is attributed to the subsampling strategy, which was employed to balance the number of male and female sentences for a more robust evaluation. Ironically, it resulted in a less stable evaluation because there is no guarantee that a randomly selected set of sentences from the larger gender set will exhibit similar bias as another randomly selected set from the same group. This issue becomes more challenging when the size of the subsampled set is significantly smaller than the original gender set. Thus, minimizing the size gap between the selected set and the original set is crucial for conducting a robust bias evaluation.

It is worth noting that the standard deviations of the SBM scores from LSG are noticeably smaller than those of Kaneko_{all} in Table 4, except for Chinese. This variance can be problematic when the score is close to the threshold of 50, potentially resulting in a contradictory finding. In this regard, MSG has an advantage over MBE and LSG as it generally discards significantly fewer sentences compared to the other two methods (Table 5) by generating new sentences for both genders, offering a more stable method for evaluating gender bias.

Table 5. The percentages of the discarded sentences for balancing both genders. MSG discards a significantly smaller number of sentences compared to the other two methods, except for German, where LSG shows the smallest amount of truncation.

Lang	MBE	LSG	MSG
English	31.98	25.67	19.60
Chinese	32.13	27.34	25.96
German	35.35	2.17	10.51
Portuguese	30.90	31.78	12.20
Spanish	31.99	32.58	15.04

While MSG provides a more consistent evaluation with a lower truncation rate compared to the other methods, it may end up producing less diverse gender words than LSG. Figure 4 shows that LSG extracts sentences with a greater number of unique gender words across the five languages than MSG. The proportion of unique gender words used to fill the [Mask] does not exceed 50% for all languages except for Chinese, where the MLM fills over 18,000 sentences using only 12 male and 4 female words. By incorporating a diverse set of vocabulary from single-gender predictions obtained using LSG, MSG can quantify the bias in MLMs even in sentences they are not inclined to generate.




Figure 4. The numbers of gender word types generated by LSG and MSG. MSG: One | Both: the number of gender word types where MSG generates for only one gender or both genders, respectively.

7. Conclusions

This paper presents robust methods for evaluating gender bias in masked language models across five languages: English, Chinese, German, Portuguese, and Spanish. Using our multilingual gender lexicon (MGL); three evaluation metrics: multilingual bias evaluation (MBE), strict bias metric (SBM), and direct comparison bias metric (DBM); and two sample generation methods: lexicon-based (LSG) and model-based (MSG), a comprehensive analysis is conducted, revealing that MSG is the most generalizable and consistent method.

As the field of bias evaluation is rapidly evolving with the emergence of new methods and metrics, we emphasize the importance of a collaborative effort from diverse perspectives to advance this research. To establish an unbiased bias evaluation system, it is essential to approach it from multiple perspectives. We hope that our work contributes to ongoing endeavors aimed at addressing gender bias and serves as an inspiration for further exploration in this critical area of research.

Author Contributions: Conceptualization, J.Y. and J.D.C.; methodology, J.Y., S.U.K., J.C. and J.D.C.; software, S.U.K. and J.Y.; validation, J.C.; formal analysis, J.Y.; investigation, J.D.C.; resources, J.Y. and J.C.; data curation, J.Y. and S.U.K.; writing—original draft preparation, J.Y., S.U.K., J.C. and J.D.C.; writing—review and editing, J.Y. and J.D.C.; visualization, J.Y.; supervision, J.D.C.; project administration, J.D.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All our resources, including datasets and evaluation scripts, are publicly available through our open source project: <https://github.com/emorynlp/GenderBiasMLM> (accessed on 6 August 2024).

Acknowledgments: The authors thank Sichang Tu and Kaustubh Dhole for their technical assistance and review. This project was made possible by the Emory NLP lab <https://www.emorynlp.org/> (accessed on 6 August 2024), and these individuals have constantly challenged us to improve our work.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

NLP	Natural Language Processing.
MLMs	Masked Language Models.
MBE	Multilingual Bias Evaluation.
AULA	All Unmasked Likelihood with Attention.
SBM	Strict Bias Metric.
LSG	Lexicon-based Sentence Generation.
MSG	Model-based Sentence Generation.
DBM	Direct Comparison Bias Metric.
MGL	Multilingual Gender Lexicon.

References

1. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.u.; Polosukhin, I. Attention is All you Need. In *Advances in Neural Information Processing Systems*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017; Volume 30.
2. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. 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 (Long and Short Papers)*, Minneapolis, MN, USA, 2–7 June 2019; pp. 4171–4186. [[CrossRef](#)]
3. Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; Stoyanov, V. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv* **2019**, arXiv:1907.11692.
4. Hartvigsen, T.; Gabriel, S.; Palangi, H.; Sap, M.; Ray, D.; Kamar, E. ToxiGen: A Large-Scale Machine-Generated Dataset for Adversarial and Implicit Hate Speech Detection. In *Proceedings of the ACL 2022, Dublin, Ireland, 22–27 May 2022*.
5. Bender, E.M.; Friedman, B. Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. *Trans. Assoc. Comput. Linguist.* **2018**, *6*, 587–604. [[CrossRef](#)]
6. Dixon, L.; Li, J.; Sorensen, J.; Thain, N.; Vasserman, L. Measuring and Mitigating Unintended Bias in Text Classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, New Orleans, LA, USA, 2–3 February 2018*; pp. 67–73.
7. Hutchinson, B.; Prabhakaran, V.; Denton, E.; Webster, K.; Zhong, Y.; Denuyl, S. Social Biases in NLP Models as Barriers for Persons with Disabilities. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020*; pp. 5491–5501. [[CrossRef](#)]
8. Clark, K.; Luong, M.T.; Le, Q.V.; Manning, C.D. ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. In *Proceedings of the International Conference on Learning Representations, Addis Ababa, Ethiopia, 26–30 April 2020*.
9. Tan, H.; Bansal, M. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. 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)*, Hong Kong, China, 3–7 November 2019; pp. 5100–5111. [[CrossRef](#)]
10. Kurita, K.; Vyas, N.; Pareek, A.; Black, A.W.; Tsvetkov, Y. Measuring Bias in Contextualized Word Representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing, Florence, Italy, 2 August 2019*; pp. 166–172. [[CrossRef](#)]

11. Zhao, J.; Wang, T.; Yatskar, M.; Ordonez, V.; Chang, K.W. 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, Volume 2 (Short Papers), New Orleans, LA, USA, 1–6 June 2018; pp. 15–20. [\[CrossRef\]](#)
12. Blodgett, S.L.; Barocas, S.; Daumé III, H.; Wallach, H. Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; pp. 5454–5476. [\[CrossRef\]](#)
13. Bender, E.M.; Gebru, T.; McMillan-Major, A.; Shmitchell, S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, New York, NY, USA, 3–10 March 2021; FAccT '21; pp. 610–623. [\[CrossRef\]](#)
14. Sun, T.; Gaut, A.; Tang, S.; Huang, Y.; ElSherief, M.; Zhao, J.; Mirza, D.; Belding, E.; Chang, K.W.; Wang, W.Y. Mitigating Gender Bias in Natural Language Processing: Literature Review. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; pp. 1630–1640. [\[CrossRef\]](#)
15. Bolukbasi, T.; Chang, K.W.; Zou, J.; Saligrama, V.; Kalai, A. Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings. In Proceedings of the NeurIPS, Barcelona, Spain, 5–10 December 2016.
16. Zhao, J.; Zhou, Y.; Li, Z.; Wang, W.; Chang, K.W. Learning Gender-Neutral Word Embeddings. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, 31 October–4 November 2018; pp. 4847–4853. [\[CrossRef\]](#)
17. Liang, S.; Dufter, P.; Schütze, H. Monolingual and Multilingual Reduction of Gender Bias in Contextualized Representations. In Proceedings of the 28th International Conference on Computational Linguistics, Barcelona, Spain, 8–13 December 2020; pp. 5082–5093. [\[CrossRef\]](#)
18. Webster, K.; Wang, X.; Tenney, I.; Beutel, A.; Pitler, E.; Pavlick, E.; Chen, J.; Petrov, S. Measuring and Reducing Gendered Correlations in Pre-trained Models. *arXiv* **2020**, arXiv:2010.06032.
19. Bommasani, R.; Davis, K.; Cardie, C. Interpreting Pretrained Contextualized Representations via Reductions to Static Embeddings. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Online, 5–10 July 2020; pp. 4758–4781. [\[CrossRef\]](#)
20. Nangia, N.; Vania, C.; Bhalerao, R.; Bowman, S.R. 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, 16–20 November 2020; pp. 1953–1967. [\[CrossRef\]](#)
21. Nadeem, M.; Bethke, A.; Reddy, S. 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), Online, 1–6 August 2021; pp. 5356–5371. [\[CrossRef\]](#)
22. Blodgett, S.L.; Lopez, G.; Olteanu, A.; Sim, R.; Wallach, H. Stereotyping Norwegian Salmon: An Inventory of Pitfalls in Fairness Benchmark Datasets. 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, 1–6 August 2021; pp. 1004–1015. [\[CrossRef\]](#)
23. Ahn, J.; Oh, A. Mitigating Language-Dependent Ethnic Bias in BERT. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Punta Cana, Dominican Republic, 7–11 November 2021; pp. 533–549. [\[CrossRef\]](#)
24. Durmus, E.; Ladhak, F.; Hashimoto, T. Spurious Correlations in Reference-Free Evaluation of Text Generation. In Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Dublin, Ireland, 22–27 May 2022; pp. 1443–1454. [\[CrossRef\]](#)
25. Kaneko, M.; Bollegala, D. *Unmasking the Mask*—Evaluating Social Biases in Masked Language Models. In Proceedings of the 36th AAAI Conference on Artificial Intelligence, Online, 22 February–1 March 2022; pp. 11954–11962.
26. Kaneko, M.; Imankulova, A.; Bollegala, D.; Okazaki, N. Gender Bias in Masked Language Models for Multiple Languages. In Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Seattle, DC, USA, 10–15 July 2022; pp. 2740–2750. [\[CrossRef\]](#)
27. Reimers, N.; Gurevych, I. Making Monolingual Sentence Embeddings Multilingual using Knowledge Distillation. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, Online, 16–20 November 2020; Association for Computational Linguistics: Stroudsburg, PA, USA, 2020.
28. Dwiastuti, M. English-Indonesian Neural Machine Translation for Spoken Language Domains. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop, Florence, Italy, 28 July–2 August 2019; pp. 309–314. [\[CrossRef\]](#)
29. Al-Haj, H.; Lavie, A. The Impact of Arabic Morphological Segmentation on Broad-coverage English-to-Arabic Statistical Machine Translation. In Proceedings of the 9th Conference of the Association for Machine Translation in the Americas: Research Papers, Denver, CO, USA, 31 October–4 November 2010.
30. Rozovskaya, A.; Roth, D. Grammar Error Correction in Morphologically Rich Languages: The Case of Russian. *Trans. Assoc. Comput. Linguist.* **2019**, *7*, 1–17. [\[CrossRef\]](#)
31. Chan, B.; Schweter, S.; Möller, T. German’s Next Language Model. In Proceedings of the 28th International Conference on Computational Linguistics, Barcelona, Spain, 8–13 December 2020; pp. 6788–6796. [\[CrossRef\]](#)

32. Cañete, J.; Chaperon, G.; Fuentes, R.; Ho, J.H.; Kang, H.; Pérez, J. Spanish Pre-Trained BERT Model and Evaluation Data. In Proceedings of the PML4DC at ICLR 2020, Virtual, 25–30 April 2020.
33. Souza, F.; Nogueira, R.; Lotufo, R. BERTimbau: Pretrained BERT Models for Brazilian Portuguese. In Proceedings of the Intelligent Systems: 9th Brazilian Conference, BRACIS 2020, Rio Grande, Brazil, 20–23 October 2020; pp. 403–417.
34. Cui, Y.; Che, W.; Liu, T.; Qin, B.; Wang, S.; Hu, G. Revisiting Pre-Trained Models for Chinese Natural Language Processing. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2020, Online, 16–20 November 2020; pp. 657–668. [[CrossRef](#)]
35. Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. Transformers: State-of-the-Art Natural Language Processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Online, 16–20 November 2020; pp. 38–45. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.