

Exploring the Impact of Training Data Distribution and Subword Tokenization on Gender Bias in Machine Translation

Bar Eluz^{◇*} Tomasz Limisiewicz^{♣*†} Gabriel Stanovsky[◇] David Mareček[♣]

[◇] School of Computer Science, The Hebrew University of Jerusalem

[♣] Faculty of Mathematics and Physics, Charles University in Prague

{bar.iluz,gabriel.stanovsky}@mail.huji.ac.il

{limisiewicz,marecek}@ufal.mff.cuni.cz

Abstract

We study the effect of tokenization on gender bias in machine translation, an aspect that has been largely overlooked in previous works. Specifically, we focus on the interactions between the frequency of gendered profession names in training data, their representation in the subword tokenizer’s vocabulary, and gender bias. We observe that female and non-stereotypical gender inflections of profession names (e.g., Spanish “doctora” for “female doctor”) tend to be split into multiple subword tokens. Our results indicate that the imbalance of gender forms in the model’s training corpus is a major factor contributing to gender bias and has a greater impact than subword splitting. We show that analyzing subword splits provides good estimates of gender-form imbalance in the training data and can be used even when the corpus is not publicly available. We also demonstrate that fine-tuning just the token embedding layer can decrease the gap in gender prediction accuracy between female and male forms without impairing the translation quality.

¹

1 Introduction

Machine translation has been one of the fastest-growing research directions in NLP. However, with the intensive growth of the technology, multiple potential harms were identified (Hovy and Spruit, 2016), including gender bias, where models rely on spurious correlations (doctors tend to be male) to make their predictions rather than more meaningful signals in their input (Stanovsky et al., 2019).

There are many reasons for the bias, such as imbalances in the training set or architecture choice. Previous works proposed various approaches to combat gender bias in translation models (Saunders

* Equal contribution.

[†] Work partially done while visiting the Hebrew University.

¹The code is available at <https://github.com/tomlimi/MT-Tokenizer-Bias>

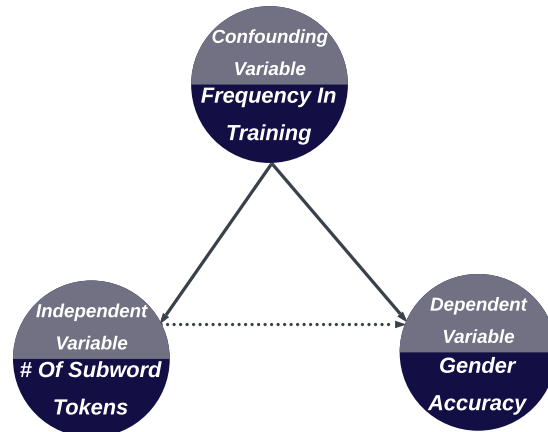


Figure 1: The schema depicts two factors affecting the accuracy of translating profession words into correct gender-inflected forms in morphologically rich languages. The first factor is the frequency of gender inflections of profession names in the training corpus, and the second is the number of subword tokens that these forms are split into. Our analysis reveals that the frequency significantly correlates with the translation accuracy and number of tokens per word. However, when we control for frequency, the correlation between the number of tokens and translation accuracy is insignificant, indicating that frequency is a confounding variable.

and Byrne, 2020; Escudé Font and Costa-jussà, 2019).

In this paper, we focus on the role of tokenization on gender bias which has been largely overlooked in previous approaches to the problem.

We want to study the causal relationship between tokenization and gender bias. Specifically, we want to know: 1. How do subword tokenizers handle different gender forms, i.e., are female and non-stereotypical gender forms split into more tokens than male and stereotypical gender forms? 2. Whether subword splitting has an impact on the accuracy of translation? 3. Would subword tokenization’s effect be significant when accounting for the frequency of gender forms in the training

corpus?

To answer those questions, we analyze pre-trained machine translation models from English to a diverse set of three languages that denote morphological gender in nouns (German, Spanish, and Hebrew).

First, we compare the number of tokens of different gender forms in the target language and find out that, indeed, female and anti-stereotypical forms are split into more tokens. Second, the causality analysis shows that the number of subword tokens may initially appear to explain the translation accuracy of gender forms. However, we find out that these factors are conditionally independent when we also consider the word frequency in the training set, as depicted in Figure 1. To support this finding, we fine-tune the model on gender balanced dataset and update its tokenizer, showing that the dataset’s role is more impactful on gender bias than tokenization.

To the best of our knowledge, this work is the first in-depth analysis of interactions between training data, tokenization method, and gender bias. Our findings confirm the previous observations (Saunders and Byrne, 2020; Zmigrod et al., 2019) indicating that the distribution of gender forms in the training data significantly influences the bias of a model. Subword tokenizers, typically trained on the same data, can also perpetuate biases present in the data. We show that it is feasible to analyze the representation of gender forms in the learned vocabulary to obtain information on gender distribution in the model’s training corpus even without having access to it.

2 Experimental Setup

2.1 Models, languages and tokenizers

The translation models we used for the analysis are OpusMT (Tiedemann and Thottingal, 2020), based on Marian NMT framework (Junczys-Dowmunt et al., 2018), trained on the Opus dataset², and mBART50 (Tang et al., 2020), a multilingual encoder-decoder model trained on 50 languages. Both OpusMT and mBART50 use SentencePiece based tokenizer (Kudo and Richardson, 2018) with unigram language model (Kudo, 2018). We test models in translating from English to German, Spanish, and Hebrew.

²<https://opus.nlpl.eu/>

Choice of target languages. We chose German, Spanish, and Hebrew as target languages since they are diverse, and all assign grammatical gender to profession names, adjectives, and nouns. Moreover, the authors of this study possess a proficiency ranging from intermediate to native levels in the aforementioned languages. To highlight typological differences between the languages: Hebrew is a Semitic language from the Afroasiatic language family with Abjad script. German is a Germanic language with the Latin alphabet. Spanish is a Romance language with Latin script that uses a change of suffix (instead of addition) for male-to-female infections. Both German and Spanish belong to the Indo-European language family.

Choice of models. We chose OPUS and mBART models because they are accessible through Huggingface, they support all the languages we selected for analysis, and they manifest strong performance in translation tasks. Both models follow state-of-the-art design choices, specifically Transformer architecture and SentencePiece tokenizer, which was shown to be preferred over BPE in multilingual models (Bostrom and Durrett, 2020a).

2.2 Data

For the evaluation of gender bias, we use the WinoMT (Stanovsky et al., 2019). It is a synthetic English dataset of sentences containing profession names coreferred by gendered pronouns. The sentences are balanced in terms of the number of male and female pronouns used. The construction of this dataset allows checking if translation models show preference toward particular gender forms. The methods for measuring these preferences are described in the following subsection 2.3.

Gender forms in target languages. To validate the gender translation correctness of professions, we collected human translations from native speakers of Hebrew, German, and Spanish (three annotators for each language) for a list of 40 professions from WinoBias dataset (Zhao et al., 2018). The list contains an equal share of professions that are majorly performed by men and women (based on labor statistics). Each profession is translated to a pair of masculine and feminine forms with the same stem (e.g., “Mediziner” – “Medizinerin” and “Arzt” – “Ärztin” are two pairs in German).³ The annotators could propose up to 3 pairs of translations

³Both pairs are German translations of the English profession “Doctor”.

for each profession. Subsequently, the authors selected a list of pairs that were proposed by at least two annotators⁴. As a result, we accepted 67 pairs for German (77% of the annotators’ propositions), 54 pairs for Hebrew (76% of the propositions), and 45 pairs for Spanish (73% of the propositions). The lists with the translations will be released upon the publication of this work.

Gender forms frequency analysis. To estimate the frequency of gender forms in the training corpus, we analyze the OPUS-100 dataset (Zhang et al., 2020). It is a sample from OPUS collection on which OpusMT was trained. It comprises multiple corpora subjects like movie subtitles, code documentation, and the Bible. We used all three training, development, and test corpora of OPUS-100, which contains 1,004,000 sentences for each of the analyzed languages.

Gender-balanced dataset for fine-tuning. We compile a gender-balanced bilingual dataset containing a simple English template: “He/She is the [profession]” paired with translations to Hebrew and German. The number of examples with male and female pronouns is equal. The templates are filled with profession names from WinoMT from English and their translations proposed by the annotators.

2.3 Metrics

Gender Translation Accuracy (F1). To evaluate the model’s performance, we check whether the model translates professions in English from WinoMT dataset to correct gender forms in the target language proposed by annotators. For instance, the profession “Physician”, when the pronouns indicate male gender in the source sentence, should be translated to corresponding male forms in German (“Arzt” or “Mediziner”). Respectively, when a pronoun indicates female gender, we should obtain female forms in the translator’s output sentence (“Ärztin” or “Medizinerin” for German). We compute the number of correct occurrences, i.e., translated sentences for which profession and gender match the English source.

We define *recall* for each profession in a specific gender as the share of correct occurrences out of a total number of source (English) sentences where the profession appeared.

⁴For detailed rules for combining the final lists and discussion of special cases, see Appendix A

The *precision* for each attested (i.e., proposed by annotators) profession translation is the share of correct occurrences out of the total number of outputted sentences where the profession translation appeared.

In our experiments, we report *F1*, which is the harmonic mean of precision and recall for each attested profession translation.

Measures of Gender Bias (ΔG , ΔS , ΔT). We use metrics proposed by Stanovsky et al. (2019) to measure gender bias in MT: ΔG measures the difference in gender translation correctness (F1) between masculine and feminine entities; similarly, ΔS measures the difference in F1 between pro-stereotypical and anti-stereotypical instances of gender role assignments.⁵

We compute ΔG for each pair of male and female translations. It is a more fine-grained approach than in previous works:

$$\Delta G = F1_{m.trans.} - F1_{f.trans.} \quad (1)$$

Analogically we compute ΔS as the difference in F1 between pro- and anti-stereotypical translation.

$$\Delta S = F1_{pro.trans.} - F1_{anti.trans.} \quad (2)$$

Additionally, we define new metrics that measure the differences in the number of tokens that distinct gender forms split into for each pair of profession name translations. ΔT_G , analogically to ΔG , quantifies the difference in the number of tokens between the male form and the female form:

$$\Delta T_G = n. \text{Tokens}_{m.trans.} - n. \text{Tokens}_{f.trans.} \quad (3)$$

While, ΔT_S corresponds to ΔS and measures the difference in number of tokens between the pro- and anti-stereotypical forms:

$$\Delta T_S = n. \text{Tokens}_{pro.trans.} - n. \text{Tokens}_{anti.trans.} \quad (4)$$

With ΔT we metricize bias already on the tokenization level and inspect its effect on the machine translation performance. We expect the words split into more tokens to be harder to predict and thus observe correlations between pairs ΔT and translation bias metrics: ΔG and ΔS .

⁵We use previous works’ attribution of stereotype to particular English professions (Zhao et al., 2018)

Examples: In case of translation from English to German, recall for English profession “Physician” in female form is the number of times “Physician” appeared in the source dataset as female and was translated to “Ärztin” or “Medizinerin”, divided by the number of times “Physician” appeared as a female in the source dataset. Precision for translation “Ärztin” is the number of times “Physician” appeared in the source dataset as female and was translated to “Ärztin” divided by the number of times the word “Ärztin” appeared in the translator’s output. F1 for translation “Ärztin” is a harmonic mean of recall for the female “Physician” and precision for “Ärztin”.

ΔG for a pair “Arzt” – “Ärztin” is difference in F1 for translation “Arzt” and translation “Ärztin”. Because “Physician” is a stereotypically male profession ΔS equals to ΔG .

Both ΔT_G and ΔT_S are the difference between the number of tokens the words “Arzt” and “Ärztin” are divided into by the German tokenizer.⁶

3 Experiments and Results

To test how tokenization affects the translation accuracy of professionals’ gender, and whether a model prefers to generate translations with fewer tokens, we design four experiments explained in the following section.

3.1 Are female and anti-stereotypical forms split into more tokens?

We take human translations (obtained from the procedure described in Section 2.2) and check how many tokens they are split into by the analyzed system’s tokenizer.

We expect that female forms will be split into more tokens than male forms, partially due to derivational suffixes appearing only in female forms.

Results Figures 2 and 3 show how many translated WinoMT professions are split into a specific number of tokens. We observe that female forms tend to be divided into more tokens than male ones. Only a small portion of female forms are not split. Similarly, pro-stereotypical translations are split into fewer tokens than anti-stereotypical ones. However, the difference is smaller than in the case of gender. We observe that the difference

in the male and female number of tokens in Spanish is smaller because female forms in Spanish are sometimes expressed by changing the suffix rather than addition (e.g. Consejero vs. Consejera).

The reason why female and anti-stereotypical forms are split into more tokens is probably that they appear in the training corpus less often.⁷

3.2 Does subword splitting affect the accuracy of translation?

We compute bias metrics: ΔG , ΔS for pairs of translations to compare them with the difference in the number of tokens between gender forms (ΔT).

We expect that when profession names differ in the number of tokens, the model will more likely generate the shorter form (typically the male or pro-stereotypical). Our inspection is that the preference for shorter forms is connected to gender bias. Therefore we expect to observe a negative correlation between the difference in translation accuracy and the number of subword tokens.

Results In Figures 4 and 5, we observe negative correlations between ΔT and both ΔG and ΔS . The relationship is stronger in the latter case. This finding supports our hypothesis that the difference in the number of tokens leads to the model’s preference for a form with fewer tokens. Additionally, for translation pairs with $\Delta T = 0$, the median of the bias measure distribution is close to zero (with the notable exception of ΔG for Hebrew). This suggests that the model is less biased for professions when both translation forms are divided into the same number of tokens.

3.3 What is the causal relationship between tokenization, training data, and gender prediction accuracy?

An alternative explanation of the negative trend observed in the previous experiment is the presence of an underlying factor, in our case, the frequency of specific gender forms in the training corpus. Previous research has shown that the terms’ frequency affects both tokenization (Kudo, 2018) and gender bias (Escudé Font and Costa-jussà, 2019).

In this experiment, we measure the significance of the correlation between those three factors. Also, we check the conditional independence between the number of tokens per target profession and gender prediction accuracy (measured by F1 score for each

⁶Please note that for stereotypically female professions, we have $\Delta G = -\Delta S$ and $\Delta T_G = -\Delta T_S$.

⁷The results of mBART50 are similar to Opus-MT and are represented in Appendix C.



Figure 2: OpusMT: Human translated profession names were grouped by the number of tokens they were split into. On the x-axis: number of tokens per word. On the y-axis: the count of male and female forms professions in each of the groups. Male forms tend to be split into fewer tokens than female forms.

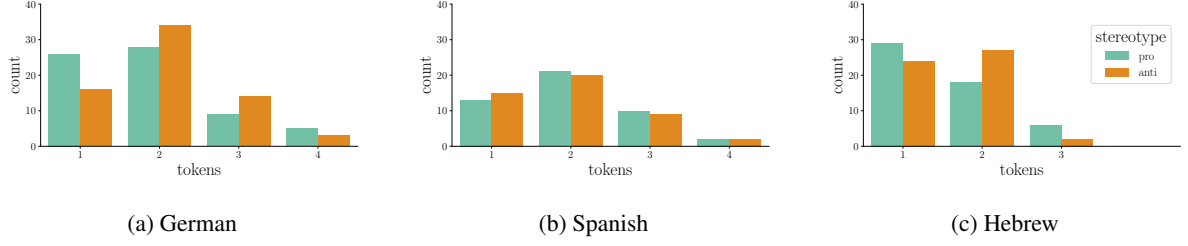


Figure 3: OpusMT: Human translated profession names were grouped by the number of tokens they were split into. On the x-axis: number of tokens per word. On the y-axis: the count of pro- and anti-stereotypical forms of professions in each of the groups. Pro-stereotypical forms tend to be split into fewer tokens than anti-stereotypical forms.

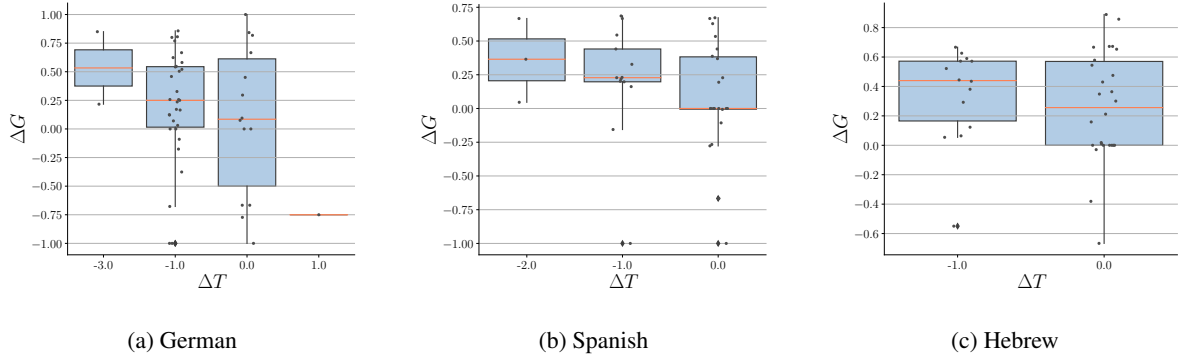


Figure 4: OpusMT: ΔG as the difference between F1-score for male and female test instances for each of paired translations. ΔT_G is the difference between the number of tokens in a male and a female form. An orange line marks the median. Pearson’s correlation coefficients and corresponding p-values: $\rho = -0.25$, $p = 0.09$ for German, $\rho = -0.21$, $p = 0.20$ for Spanish, and $\rho = -0.11$, $p = 0.50$ for Hebrew.

target profession described in Section 2) given the profession form’s frequency in the training corpus:

$$F1 \perp n. \text{ Tokens} \mid \text{Frequency} \quad (5)$$

Results In Figure 6, we observe that as the number of tokens decreases, F1 increases. Taking into account the correlation coefficient, the F1 measure is more sensitive to the frequency of a word in the training corpus. Moreover, less frequent words tend to be split into more tokens. From the density plots

(on the diagonal of the figure), we see that male profession words (especially pro-stereotypical ones) appeared much more often in the training corpora. Thus, they tended to be split into fewer tokens.

All the correlations between frequency and the two remaining factors are statistically significant ($p < 0.05$), while the correlation between the number of tokens and F1 score is significant only for German.

We performed Jonckheere-Terpstra test (Jonckheere, 1954) to check conditional independence as described in Equation 3.3. The test showed that

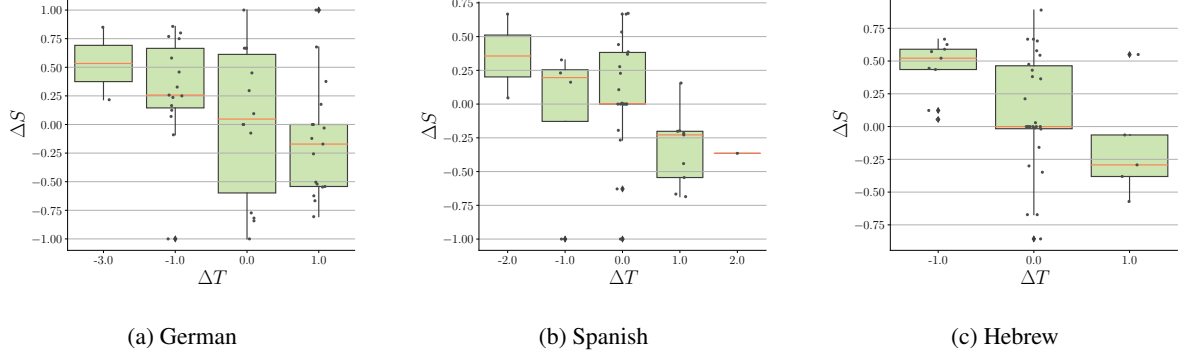


Figure 5: OpusMT: ΔS as the difference between F1-score for pro- and anti-stereotypical test instances for each of paired translations. ΔT_G is the difference between the number of tokens in a pro- and an anti-stereotypical translation. An orange line marks the median. We observe a significant negative correlation between the two measures. Pearson’s correlation coefficients and corresponding p-values: $\rho = -0.37$, $p = 0.010$ for German, $\rho = -0.37$, $p = 0.024$ for Spanish, and $\rho = -0.41$, $p = 0.008$ for Hebrew.

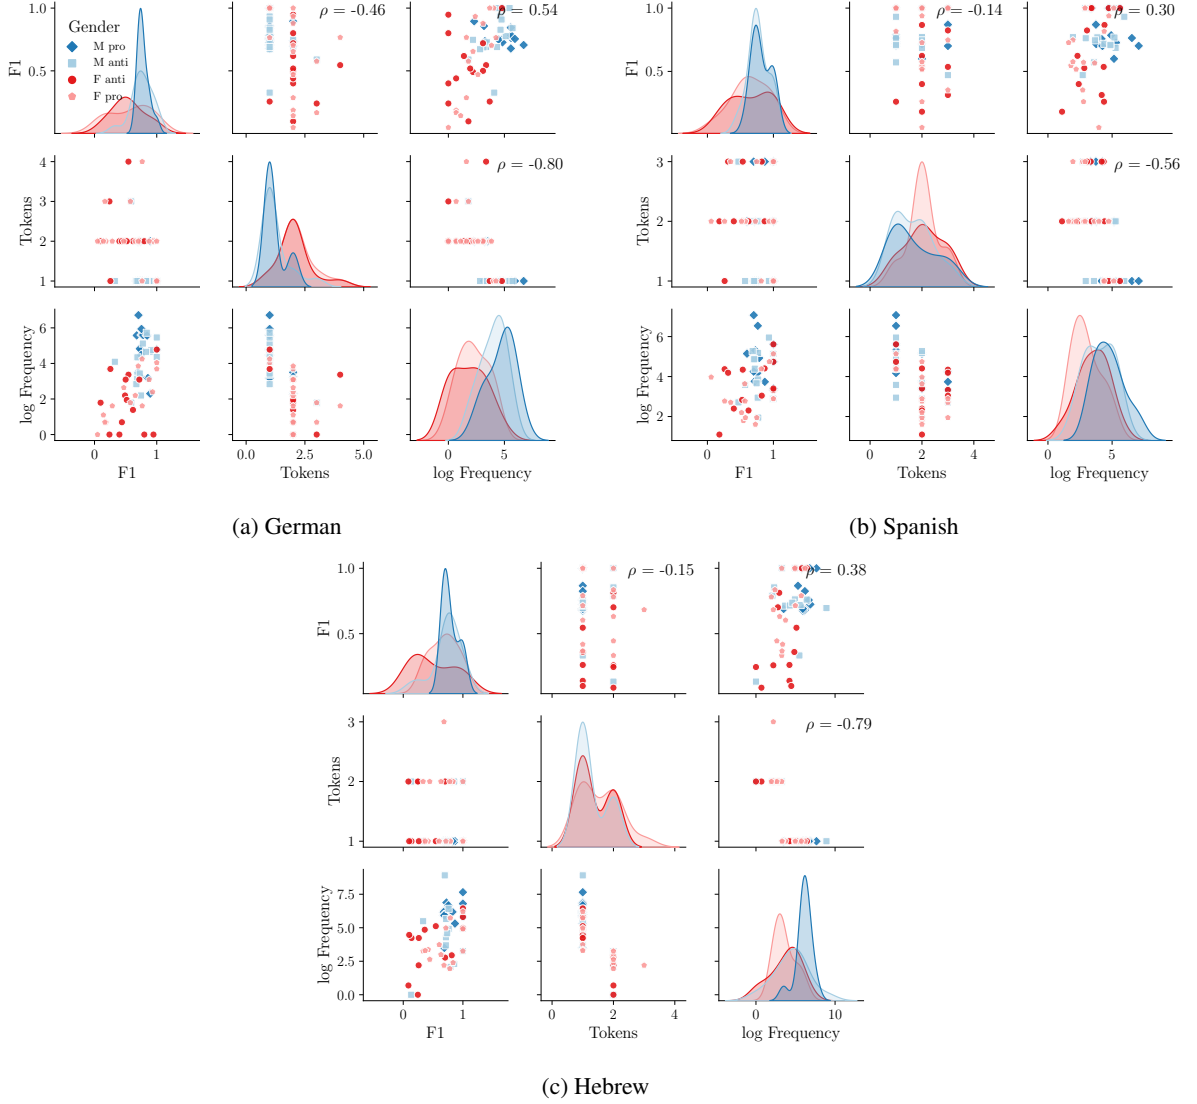


Figure 6: Pair analysis of gender prediction performance (F1), number of tokens, and the frequency of each profession in the OPUS-100 dataset. Each style of dots represents professions in the male/female pro-stereotypical/anti-stereotypical form. The diagonal plots show the density of the feature for the specific gender and stereotype sets.

conditional independence cannot be rejected for all the target languages ($p = 0.78$ for German, $p = 0.62$ for Spanish, and $p = 0.39$ for Hebrew). Those results show that the frequency of a word is a confounding factor affecting both the number of tokens per profession word and F1 scores. **Hence subword splitting of profession words is not a significant contributor to the correctness of gender prediction given the frequency of gender forms in the training corpus.**

3.4 Will intervening in the model’s training data and the tokenizer’s vocabulary reduce gender bias?

To verify our findings about causes of gender bias, we propose two interventions in the translation models to German, Spanish, and Hebrew: 1. fine-tuning on a dataset with an equal number of male and female forms;⁸ 2. Adapting the tokenizer’s vocabulary by adding all translations proposed by annotators to assert that they will not be split into subword tokens.

We monitor standard gender bias measures from Stanovsky et al. (2019): gender accuracy, ΔG , and ΔS ,⁹ and also BLEU on OPUS-100 test split to check if fine-tuning leads to deterioration of translation quality. To determine if the potential improvement results from embeddings’ fine-tuning or adding profession words to the vocabulary, we evaluate the baseline, where the embedding layer is fine-tuned without updating the tokenizer’s vocabulary.

Results Table 1 shows that fine-tuning embedding layers improve the accuracy of translating to the correct gender and decrease preference of male forms (ΔG), while the quality of translation stays on a similar level as in the original model. Protecting gender forms from splitting (by adding them to the tokenizer’s vocabulary) only slightly further converges ΔG towards zero for German while bringing over 2 points drop in BLEU. Performing the vocabulary update before fine-tuning models deteriorates bias measures and BLEU for Hebrew. For Spanish, ΔS improves while ΔG and BLEU worsen.

⁸The dataset contains the sentences in the form “He/She is the [profession]” translated to the target language and filled with professions from WinoMT.

⁹In this experiment, we use the original formulation of ΔG and ΔS . It means $F1$ is computed for the whole dataset, instead of particular professions as in the previous experiments.

Interestingly, fine-tuning the embedding layer increases the stereotypical bias (ΔS) for German and Spanish and decreases it only slightly for Hebrew, suggesting that this method could not be sufficient for mitigating this type of bias.

The results confirm our previous observation that the training (or fine-tuning) dataset has more impact on gender bias encoded in the model than subword splitting.

4 Related Work

Evaluation and Mitigation of Gender Bias in NMT Escudé Font and Costa-jussà (2019) created a set of sentences of the format “I’ve known [him/her/Mary/John] for a long time, my friend works as a [profession].” They translated all sentences to Spanish and checked whether “friend” was translated as “amigo” (male) or “amiga” (female). They found out that “Him” is predicted at almost 100% accuracy for all models, but the accuracy drops when predicting the word “her” on all models.

Stanovsky et al. (2019) created a mechanism to evaluate gender bias in machine translation. They showed that almost all translation systems perform significantly better on male roles.

The past works acknowledged the effect of gender form imbalance in training corpus on gender bias manifested in the models. Specifically, Zmigrod et al. (2019) reduce the bias by training the NMT systems on data augmented by adding female forms. Saunders and Byrne (2020); Costa-jussà and de Jorge (2020) propose a debiasing algorithm based on fine-tuning the model on a dataset with a comparable number of male and female forms. These approaches are more sustainable because they do not require training the model from scratch. These approaches are in line with our observation that the imbalance of gender forms in training data is the key source of gender bias in the model.

Savoldi et al. (2021) survey the methods for evaluating and mitigating bias in machine translation, identify risks connected to them and propose direction for their improvement.

Guo et al. (2022) propose an automatic prompt-based method to mitigate the biases in pretrained language models. They identify biased prompts and propose a distribution alignment loss to mitigate it.

The Role of Tokenization on System’s Performance Domingo et al. (2018) show that tokeniz-

	Accuracy			ΔG			ΔS			BLEU		
	DE	ES	HE	DE	ES	HE	DE	ES	HE	DE	ES	HE
original	62.2	56.3	52.1	9.6	16.4	12.3	12.5	16.7	40.4	31.4	42.8	36.0
FT	67.1	58.9	54.0	-2.9	10.9	7.8	21.5	20.0	35.1	31.3	42.0	35.8
FT + VU	67.8	58.0	52.4	1.1	14.0	10.4	21.7	18.2	40.7	29.1	38.3	35.1

Table 1: Embedding fine-tuning results in German and Hebrew. Original OpusMT model results compared with the model after embedding layer fine-tuning (row 2) and the model after embedding layer fine-tuning and updating vocabulary with all profession gender forms in target language (row 3).

ers play a significant role in the neural machine translation pipeline. The tokenization of the target language affects evaluation measures (BLEU) by up to 20 points.

Bostrom and Durrett (2020b) compare the popular subword tokenizers: Unigram (Kudo, 2018) and BPE (Sennrich et al., 2016). They show that the former more often splits words on morphological boundaries and thus can improve the model’s performance on downstream tasks.

While the recent survey (Mielke et al., 2021) shows that the choice of effective tokenization method depends on the task, and there is no specific tokenization algorithm suiting all applications.

The role of tokenization on gender bias was not widely evaluated in past works. Libovický et al. (2022) showed that character-based translators to morphologically rich languages (Czech and German) obtain similar bias results as subword-based systems, even though female forms in these languages contain relatively more characters. It aligns with our observation that tokenization is not a significant source of bias.

Gaido et al. (2021) explored how segmenting methods influence gender bias in speech translation models. Such models have different features, such as vocal characteristics used to measure gender bias over all words, while our work focuses on gender bias in profession words. In contrast to their work, we focus on formal gender bias definitions and analyzing the relation between frequency, accuracy, and number of tokens.

5 Discussion and Future Work

Our confirms the validity of the causal schema depicted in Figure 1 explaining the lower translation accuracy for female and anti-stereotypical profession names. We consciously analyzed the dataset of only non-ambiguous sentences where the gender of each profession is known. This selection was

made to enable us to evaluate the accuracy of the ground-truth gender.

Furthermore, the stereotypical occupations are based on US Department of Labor statistics, but it cannot be guaranteed that these same stereotypes are present in other cultures. However, it can be considered a reasonable estimate for other languages, as evidenced by the observation that non-stereotypical professions appear less frequently in their training corpora. Future research may analyze gender roles in target languages to corroborate these observations.

Another future point of interest is mapping more factors for bias by isolating more features. The dependencies of tokens, frequency, and gender bias will be examined on a larger scale, with more words and different types of tokenization. We also intend to broaden the scope of the analysis to other languages and include neutral gender forms for the already analyzed languages.

6 Conclusions

Our study found that profession words in German, Hebrew, and Spanish tend to be split into more tokens in the female form than in the male form. We then investigated whether this phenomenon amplifies the tendency of NMT models to translate male professions more accurately than female ones. Our results showed that the frequency of gender forms confounds the relationship between the number of tokens and gender bias in the training set. However, the number of subword tokens per word can be used to estimate its frequency in the unavailable training corpus.

The findings of our analysis of translation models were supported by the trends observed in the results of the proposed debiasing method. Specifically, we found that fine-tuning token embeddings on a gender-balanced dataset had a more significant impact on reducing bias than updating the

tokenizer’s vocabulary with underrepresented gender forms.

These findings suggest that future research should also focus on other aspects of NMT models in order to mitigate gender bias effectively.

Acknowledgments

We thank anonymous reviewers for their valuable comments on the previous versions of this article. This work was partially supported by the Israeli Ministry of Science and Technology (grant no. 2088) and by grant 23-06912S of the Czech Science Foundation. Tomasz Limisiewicz’s visit to the Hebrew University has been supported by grant 338521 of the Charles University Grant Agency and the Mobility Fund of Charles University. We have been using language resources and tools developed, stored, and distributed by the LINDAT/CLARIAH-CZ project of the Ministry of Education, Youth and Sports of the Czech Republic (project LM2018101).

Limitations

As pointed out in the last section, tokenization can reinforce gender bias in the models but is not the sole factor responsible for gender bias. We expect that even if male and female forms were tokenized to the same number of tokens, bias could still persist in the models.

The important aspect is the frequency of words in the training corpus, which also affects tokenization. We did not have exact information on the model’s training corpus. However, the analysis of the OPUS-100 dataset, which was sampled from the full dataset, showed that male forms are much more frequent than female ones.

In addition, we analyzed only 40 stereotypical professions from WinoMT, which is only a subset of the whole dataset. Extending the analysis to other professions without typical gender-role assigned can reveal slightly different trends.

Finally, non-zero bias measures indicate the presence of bias in the model, but the opposite implication is not true. Even in the case where both ΔG and ΔS are close to zero, we cannot claim that the model is unbiased, as there may be other bias manifestations of bias involved.

A significant limitation of this work is that we compare only male and female forms of the professions while neglecting neutral forms. We expect neutral forms to be even harder for the model to

predict than female forms because of the added suffixes that are rarely present in the training data. For instance, German uses forms with “*” and female prefixes (e.g., “Mechaniker*in”) while “\” is used in Hebrew for that purpose. Our initial analysis shows that neutral forms are sporadic in the training corpora. For instance, in German, only 7 neutral forms are using the inclusive suffix “*in”, compared to 1351 female forms of the analyzed professions. The proportion of neutral:female:male forms is approximately 1:193:1845.

Ethics Statement

We introduce new methods for evaluating and mitigating gender bias in machine translation. We believe that this solution will incentivize the development of fairer models in the future.

We do not identify any ethical risks connected to this work.

References

- Kaj Bostrom and Greg Durrett. 2020a. Byte pair encoding is suboptimal for language model pretraining. *arXiv preprint arXiv:2004.03720*.
- Kaj Bostrom and Greg Durrett. 2020b. [Byte pair encoding is suboptimal for language model pretraining](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4617–4624, Online. Association for Computational Linguistics.
- Marta R. Costa-jussà and Adrià de Jorge. 2020. [Fine-tuning neural machine translation on gender-balanced datasets](#). In *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*, pages 26–34, Barcelona, Spain (Online). Association for Computational Linguistics.
- Miguel Domingo, Mercedes García-Martínez, Alexandre Helle, Francisco Casacuberta, and Manuel Heranz. 2018. How much does tokenization affect neural machine translation? *ArXiv*, abs/1812.08621.
- Joel Escudé Font and Marta R. Costa-jussà. 2019. [Equalizing gender bias in neural machine translation with word embeddings techniques](#). In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 147–154, Florence, Italy. Association for Computational Linguistics.
- Marco Gaido, Beatrice Savoldi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. How to split: the effect of word segmentation on gender bias in speech translation. *arXiv preprint arXiv:2105.13782*.
- Yue Guo, Yi Yang, and Ahmed Abbasi. 2022. [Auto-debias: Debiasing masked language models with automated biased prompts](#). In *Proceedings of the*

- 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1012–1023, Dublin, Ireland. Association for Computational Linguistics.
- Dirk Hovy and Shannon L. Spruit. 2016. [The social impact of natural language processing](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 591–598, Berlin, Germany. Association for Computational Linguistics.
- A. R. Jonckheere. 1954. A distribution-free k-sample test against ordered alternatives. *Biometrika*, 41:133–145.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast neural machine translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Jindřich Libovický, Helmut Schmid, and Alexander Fraser. 2022. [Why don’t people use character-level machine translation?](#) In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2470–2485, Dublin, Ireland. Association for Computational Linguistics.
- Sabrina J. Mielke, Zaid Alyafeai, Elizabeth Salesky, Colin Raffel, Manan Dey, Matthias Gallé, Arun Raja, Chenglei Si, Wilson Y. Lee, Benoît Sagot, and Samson Tan. 2021. [Between words and characters: A brief history of open-vocabulary modeling and tokenization in nlp](#).
- Danielle Saunders and Bill Byrne. 2020. [Reducing gender bias in neural machine translation as a domain adaptation problem](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7724–7736, Online. Association for Computational Linguistics.
- Beatrice Savoldi, Marco Gaido, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. [Gender Bias in Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 9:845–874.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2019. [Evaluating gender bias in machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1679–1684, Florence, Italy. Association for Computational Linguistics.
- Yuqing Tang, Chau Tran, Xian Li, Peng-Jen Chen, Naman Goyal, Vishrav Chaudhary, Jiatao Gu, and Angela Fan. 2020. Multilingual translation with extensible multilingual pretraining and finetuning. *arXiv preprint arXiv:2008.00401*.
- Jörg Tiedemann and Santhosh Thottingal. 2020. OPUS-MT — Building open translation services for the World. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation (EAMT)*, Lisbon, Portugal.
- Biao Zhang, Philip Williams, Ivan Titov, and Rico Sennrich. 2020. [Improving massively multilingual neural machine translation and zero-shot translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1628–1639, Online. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. [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)*, pages 15–20, New Orleans, Louisiana. Association for Computational Linguistics.
- Ran Zmigrod, Sabrina J. Mielke, Hanna Wallach, and Ryan Cotterell. 2019. [Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1651–1661, Florence, Italy. Association for Computational Linguistics.

A Human Translation

A.1 Post-processing of Collected Data

The post-processing of the human translations was done according to the following rules: First, we kept the translations suggested by at least two translators. If there's no translation option suggested at least two times, we decided which translations to keep according to our knowledge of the language. This happened only for the following translations in Hebrew: "attendant", "construction worker", and "analyst". The translation selection for those words was made by an author who is a native Hebrew speaker. For professions where translators propose the same word as both male and female translations, we keep those as valid pairs.

In case spelling is inconsistent among annotators, we kept the more common spelling while fixing pronunciation mistakes according to our knowledge. Lastly, we keep at most five pairs of translations per profession.

A.2 Instructions for Annotators

The goal of the task is to get the correct translation of profession words into your language for both Male and Female forms.

1. Select the language you want to translate to (your native language: German, Spanish, or Hebrew).
2. For each word in English, write up to 3 translations both in Male and Female form.
3. Organize the translation into pairs that differ only in endings (e.g., words: Berater and Beraterin are a pair but Beraterin and Ratgeber are not a pair in German). Put these words in the neighboring cells ("Translation N Male" and "Translation N Female").
4. in case there is no translation in the other gender that differ only in an ending, leave the other cell of the pair blank.
5. If you propose less than 3 translations in each gender, leave the last cells in a row empty.

B Details of Fine-tuning

In fine-tuning, we train only the translation model's input/output embedding layer while keeping the rest of the parameters frozen. We further trained the model for three epochs on the gender-balanced

dataset (Section 2.2) with batches consisting of 16 examples. For optimization, we used Adam (Kingma and Ba, 2015) with default parameters with a constant learning rate of $5e - 5$. The fine-tuning took about 10 minutes on Nvidia A40 GPU.

C MBART results

We repeat the analysis of the number of tokens per word described in Section 3.1 for MBART. Figures 7 and 8 show that similarly to Opus models, MBART split female and anti-stereotypical words in the target language into more tokens. Noticeably, almost no female words, in all the languages are represented as one token.

Figures 9 and 10 show similar trends for ΔG and ΔS against ΔT as observed in Section 3.2 for OPUS models.

D WinoMT details

The WinoMT dataset that we used for evaluation contains 3,888 sentences. Each sentence contains a profession and a pronoun that describes the gender of the profession. The dataset is equally balanced between male and female examples and also between stereotypical and non-stereotypical gender-role assignments. Specifically, there are 1826 sentences with a male pronoun assigned to the profession and 1822 sentences with a female pronoun. The remaining 240 sentences are gender-neutral, i.e., they contain the pronoun "they".

Note that this dataset contains only non-ambiguous sentences where the gender of each profession is known. This selection was made to enable us to evaluate the accuracy of the ground-truth gender.

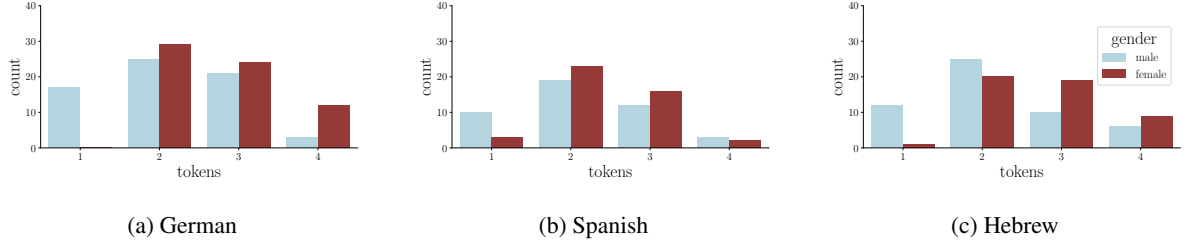


Figure 7: MBART: Human translated profession names grouped by the number of tokens they were split into. On the x-axis: number of tokens per word. On the y-axis: the count of male and female forms professions in each of the groups. Male forms tend to be split into fewer tokens than female forms.

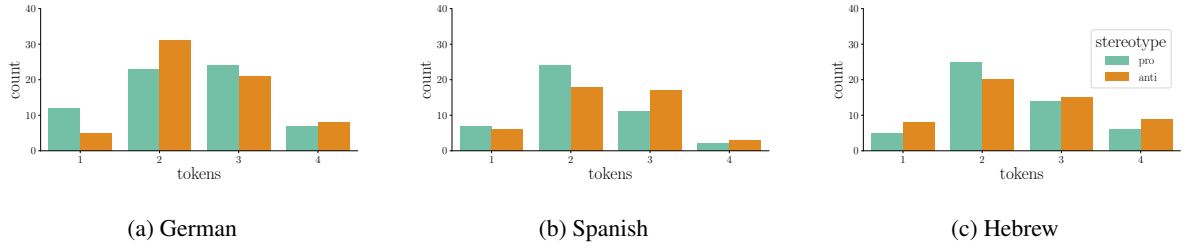


Figure 8: MBART: Human translated profession names grouped by the number of tokens they were split into. On the x-axis: number of tokens per word. On the y-axis: the count of pro- and anti-stereotypical forms of professions in each of the groups. Pro-stereotypical forms tend to be split into fewer tokens than anti-stereotypical forms.

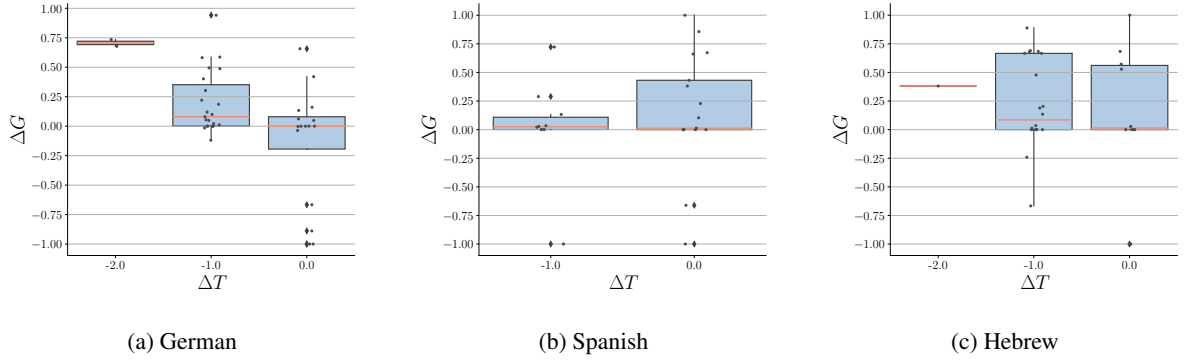


Figure 9: MBART: ΔG as the difference between F1-score for male and female test instances for each of paired translations. ΔT is the difference between the number of tokens in a male and a female form. The median is marked by an orange line.

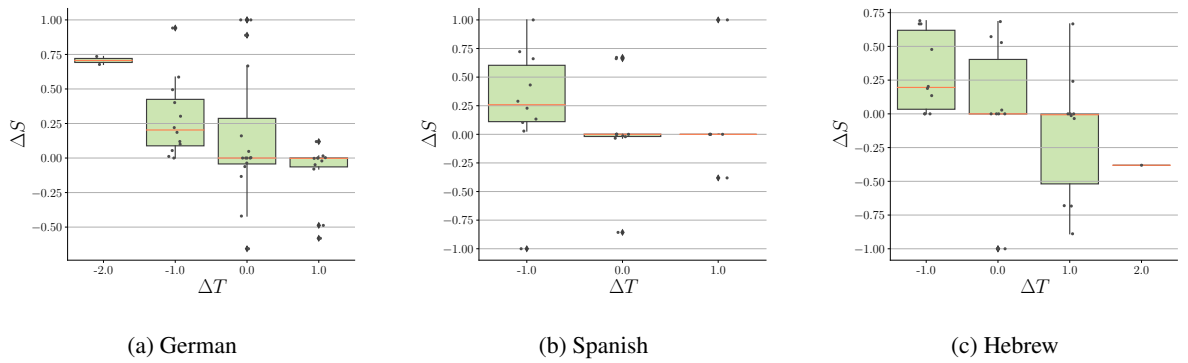


Figure 10: MBART: ΔS as the difference between F1-score for pro- and anti-stereotypical test instances for each of paired translations. ΔT is the difference between number of tokens in a pro- and an anti-stereotypical translation. Median is marked by an orange line.