

# Evaluating Gender Bias in Speech Translation

Marta R. Costa-jussà<sup>1</sup>, Christine Basta<sup>1,2</sup>, Gerard I. Gallego<sup>1</sup>

<sup>1</sup> TALP Research Center, Universitat Politècnica de Catalunya, Barcelona, Spain

<sup>2</sup> Institute of Graduate Studies and Research, Alexandria University, Egypt

{marta.ruiz, christine.raouf.saad.basta, gerard.ion.gallego}@upc.edu

## Abstract

The scientific community is increasingly aware of the necessity to embrace pluralism and consistently represent major and minor social groups. Currently, there are no standard evaluation techniques for different types of biases. Accordingly, there is an urgent need to provide evaluation sets and protocols to measure existing biases in our automatic systems. Evaluating the biases should be an essential step towards mitigating them in the systems. This paper introduces WinoST, a new freely available challenge set for evaluating gender bias in speech translation. WinoST is the speech version of WinoMT, an MT challenge set, and both follow an evaluation protocol to measure gender accuracy. Using an S-Transformer end-to-end speech translation system, we report the gender bias evaluation on four language pairs, and we reveal the inaccuracies in translations generating gender-stereotyped translations.

**Keywords:** challenge set, multilinguality, direct speech-to-text translation

## 1. Introduction

There is a massive lack of representation of diverse gender, race, cultural groups in the power structures. These problems permeate data science causing unbalanced progress in this area, which misrepresents social groups by gender, race, and nationality.

Biases have been shown in Machine Translation (MT) systems when translating from high-resource languages to low-resourced ones (Aharoni et al., 2019) and from less-morphological languages to high-morphological ones producing gender-stereotyped forms in inaccurate translations (Stanovsky et al., 2019; Prates et al., 2020; Saunders and Byrne, 2020a). Additionally, Automatic Speech Recognition (ASR) has demonstrated biases having a higher error rate for female voices than for males (Tatman, 2017). Speech Translation (ST) is at the intersection of ASR and MT tasks, perpetuating biases from both tasks.

While this is still an uncovered challenge in most artificial intelligence tasks, the scientific community is making huge efforts towards bringing to light this challenge. In the short term, researchers are devoting efforts towards bias evaluation, detection, and mitigation methods. On the other hand, there is multidisciplinary work in education, politics, and communications in the long-term perspective.

For addressing this problem in MT, some papers propose evaluating gender bias in MT (Stanovsky et al., 2019; Levy et al., 2021), creating multilingual balanced data to train fairer systems (Costa-jussà et al., 2020), or using equalizing techniques to mitigate the effect of unbalanced training data (Font and Costa-jussà, 2019; Saunders and Byrne, 2020a). Finally, in ST, (Bentivogli et al., 2020) have proposed the MuST-SHE corpus, a natural benchmark for automatically evaluating gender bias.

The main contribution of our work is providing a large-

scale multilingual ST challenge set,<sup>1</sup> which follows an evaluation protocol for the analysis of gender bias, previously proposed for MT (Stanovsky et al., 2019) (see Table 1). This set mainly evaluates the gender bias by analysing the inflection of the gendered words in the sentences, from English to any other language. Additional bias can be detected considering different speakers' gender, but such bias is out of the scope of our challenge set. Although the synthetic challenge set can propagate stereotypes as described by (Blodgett et al., 2020), it still remains valuable to evaluate and monitor the model's performance and provide insights into how the model treats gender-related issues (Blodgett et al., 2021).

## 2. Bias Statement

As proposed in previous work (Blodgett et al., 2020) and suggested in related venues,<sup>2</sup> we formulate the bias statement of our work. Our work focuses on proposing a challenge set and benefit from an evaluation protocol for the application of ST. Our challenge set serves to measure how accurately gendered source spoken words are translated. It includes stereotypical and anti-stereotypical examples, adding value to the evaluation analysis.

The main contribution of our proposal is to have an objective evaluation protocol that can measure how biased our ST systems are. We consider that our work is based on a synthetic dataset simplified to a binary gender, which does not aim to be a mirror of reality. However, we are contributing to exhibiting that current systems contain biases while sustaining that long-term policies are required to solve this problem (D'Ignazio and Klein, 2018).

<sup>1</sup>Freely available in Zenodo (10.5281/zenodo.4139080)

<sup>2</sup><https://genderbiasnlp.talp.cat/gebnlp2020/how-to-write-a-bias-statement/>

Table 1: ST Gender Bias (GB) benchmarks comparison, MuST-SHE and WinoST with respect to the GB in speaker’s voice, GB in sentence content and language pairs covered by each benchmark. X stands for any language, meaning that WinoST can be evaluated from English to any other language.

	MuST-SHE	WinoST (Ours)
<i>GB in sent. content</i>	✓	✓
<i>GB in speaker voice</i>	✓	
<i>Language pairs</i>	En → Fr, Es, It	En → X

### 3. Gender Bias within MT systems and Related Work

Grammatical gendered languages have richer grammar for expressing gender. In these languages, gender has to be assigned to all nouns, and consequently, all the articles, verbs, and adjectives have to correspond with the gender of this noun. This leads to incorrect translation from a natural gendered language (English) to a grammatical gendered marking language (Spanish) due to the lack of explicit evidence of the gender in the source (Gonen and Webster, 2020). Thus, gender bias occurs due to the gender information loss and over-prevalence of gendered forms in the training data. An illustrative example would be the translation from *The doctor spoke to the responsible nurse in the hospital.* to its Spanish version. The Spanish translation would be *El médico habló con la enfermera responsable en el hospital.*, assuming the nurse to be female and the doctor to be male. Although nothing in the text mentions their genders, the systems would translate in a biased manner. Generally speaking, MT systems have proven to have biased outputs. Even in more context, translations seem to ignore the context in sentences with the masculine and stereotyped gendered roles (Stanovsky et al., 2019; Saunders and Byrne, 2020a). The main reason for such bias is training models on human-biased data (Costa-jussà, 2019).

Researchers have recently dedicated efforts to attempt resolving such bias. We can thereby describe three research lines of gender bias in MT: mitigating gender bias in MT, detecting issues in translation, and creating the challenge test sets to evaluate gender bias in the systems.

Regarding the first research line, approaches have been dedicated to solving the translation bias from a neutral gendered language to a grammatically gendered marking one. Adding a gender tag of the speaker during training enhances the translation quality, as demonstrated in (Vanmassenhove et al., 2018). It facilitates the gender prediction correctly when translating from English to other gendered languages, giving control over the translation hypothesis gender. This was confirmed in recent work by (Basta et al., 2020), who proved that adding gender helps to increase the accuracy of gendered translations. Moreover, the authors showed that increasing context has a better effect on gendered translations, leading to higher per-

formance. (Moryossef et al., 2019) incorporate gender information by prepending a short phrase for each sentence in inference time, which acts as a gender label. Recent work (Saunders and Byrne, 2020b; Costa-jussà and de Jorge, 2020) has treated gender biasing as a domain-adaptation problem, in which the system is fine-tuned instead of retrained for mitigating gender bias. They have adapted a set of synthetic sentences with equal numbers of entities using masculine and feminine inflections to fine-tune the MT system. (Font and Costa-jussà, 2019) introduced the idea of adjusting the word embeddings, which improved performance on an English-Spanish occupations task.

Concerning research for detecting gender translation issues, (Prates et al., 2020) have examined Google translations and proved that mentions of stereotyped professions are more reliably translated than those anti-stereotyped. In their study, they have used sentence templates filled with word lists of professions and adjectives. The authors in (Cho et al., 2019) also have studied the pronouns’ translations for the Korean language using sentence templates. The recent work (Gonen and Webster, 2020) has proposed a BERT-based perpetuation method to identify gender issues in MT automatically. The technique discovers new sources of bias beyond the word lists used previously.

The most related line of research to our work is creating challenge sets. The first MT challenge set was introduced by (Stanovsky et al., 2019), called WinoMT. WinoMT is a test set of 3,888 sentences, where each sentence contains two human entities, where one of them is co-referent to a pronoun. The evaluation depends on comparing the translated entity with the golden gender, with the objective of a correctly gendered translation. The authors identified three metrics for the evaluation: accuracy,  $\Delta G$  and  $\Delta S$ . The accuracy is the percentage of correctly gendered translated entities compared to the gender of golden entities.  $\Delta G$  is the difference in F1 score between the set of sentences with male entities and female entities set.  $\Delta S$  is the difference in accuracy between the set of sentences with pro-stereotypical entities and the set with anti-stereotypical entities. A pro-stereotypical set identifies ‘developer’ as male and ‘hairdresser’ as female. An anti-stereotypical identifies the former as female and the latter as male. As far as we know, the only existing related work for studying gender bias in speech is presented in (Bentivogli et al., 2020). They created a

benchmark dataset that evaluates two categories of gender bias: one related to the speaker gender and another to the utterance content. The benchmark is limited to English-to-French, English-to-Spanish and English-to-Italian language pairs.

Given that the closest work to ours is MuST-SHE (Bentivogli et al., 2020), as follows we detail the key difference between the two resources: while MuST-SHE contains naturally occurring gender phenomena, WinoST is a synthetic challenge set. This difference has several implications, not only in terms of size of the resources themselves (generating synthetic data is somehow easier than collecting them in the wild) but also in terms of their applicability in realistic evaluation settings. Both types of resources are useful and much needed. Therefore, we see some complementarity here rather than concurrent alternatives where one is better than the other.

#### 4. Speech Translation System

We trained a ST system to evaluate its gender bias with the methodology we are presenting. We used an end-to-end ST approach that directly translates the utterance without obtaining the intermediate transcriptions. This task was introduced by (Bérard et al., 2016), and recently it had a growing interest in the research community (Weiss et al., 2017; Vila et al., 2018; Liu et al., 2019). The data we used to train it is the MuST-C corpus, that consists of speech fragments from TED Talks, its transcriptions and translations into 8 European languages (Di Gangi et al., 2019a).

The architecture we used is the S-Transformer, a popular adaptation of the Transformer for ST (Di Gangi et al., 2019b). It applies a stack of convolutions and self-attention layers to process the log-Mel spectrograms extracted from the speech utterances. The two bidimensional (2D) convolutional layers are in charge of capturing local patterns in the spectrogram, in both time and frequency dimensions. Moreover, they reduce the features maps by four, which is crucial to avoid memory issues that happen when feeding the Transformer with long sequences. Then, the two 2D self-attention layers, which were introduced by (Dong et al., 2018), model long-range dependencies that convolutional layers cannot capture. Finally, the self-attention layers of the Transformer encoder also include a logarithmic distance penalty that biases them towards the local context (Sperber et al., 2018).

Following a common approach, we pre-trained the S-Transformer encoder for ASR to improve the performance of the final ST system, as introduced by (Bérard et al., 2018) and recommended by the authors of the S-Transformer.

#### 5. Proposed Gender Evaluation: WinoST challenge set

WinoST is the speech version of WinoMT, recorded in off-voice by an American female speaker, and consists




Figure 1: WinoST Evaluation Block Diagram for Speech Translation

of 3,888 speech audios in English. By nature, sentences from WinoST contain information in the utterance content, not in gender information in the speaker’s voice. An example of these sentences is *The developer argued with the designer because she did not like the design.*, where *she* refers to *developer*, meaning that the *developer* is actually a female.

WinoST serves as an input of the ST system to be evaluated, and the output text of the systems follows the same evaluation protocol as WinoMT (Stanovsky et al., 2019). Figure 1 shows the block diagram of this procedure. As a side-product, and not shown in the figure, WinoST can also be used as a challenge set for evaluating ASR gender bias.

Further technical details on WinoST are reported in Table 2, including number of files, total hours/words, audio recording and format. The voice mastering process we applied to the recordings includes dynamic voice processing, broadcasting, equalization and filtering. WinoST is available under the MIT License<sup>3</sup> with the limitation that recordings cannot be used for speech synthesis, text to speech, voice conversion or other applications where the speaker voice is imitated or reproduced.

Table 2: WinoST details.

<b># Files</b>	3,888
<b># Hours</b>	~ 6
<b># Words</b>	~ 50,500
<b>Audio format</b>	WAV (48 KHz, 16-bit)

#### 6. Experiments

In this section, we are describing the first experiments with WinoST. We limit our experiments to four language pairs, but WinoST is extendable to any language pair with English as source language. The only requirement is to have a part-of-speech for the target language. We describe the baseline ST systems that we are using and the results that we obtain in gender accuracy.

<sup>3</sup>[https://github.com/gabrielStanovsky/mt\\_gender/blob/master/LICENSE](https://github.com/gabrielStanovsky/mt_gender/blob/master/LICENSE)

## 6.1. Data preprocessing

Before training the model, we preprocessed both speech and text data. We extracted 40-dimensional log-Mel spectrograms from the audio files, using a window size of 25 ms and hop length of 10 ms, with XNMT (Neubig et al., 2018).<sup>4</sup> We normalized the punctuation from text data, we tokenized it, and we de-escaped special characters, using the Moses scripts.<sup>5</sup> Furthermore, in the case of transcriptions, we lowercased them, and we removed the punctuation. We used the BPE algorithm (Sennrich et al., 2016) for encoding translation texts, using a vocabulary size of 8000 for each language, but a character-level encoding in the case of transcriptions.

## 6.2. System Details

The model we used has two convolutional layers with a kernel size of 3, 64 channels and a stride of 2. The Transformer has an embedding size of 512, 6 layers at the encoder and decoder, 8 self-attention heads, and a feed-forward network hidden size of 1024. We trained the S-Transformer with an Adam optimizer, with a learning rate of  $5 \cdot 10^{-3}$ , and an inverse square root scheduler. The training has a warm-up stage of 4000 updates, in which the learning rate grows from  $3 \cdot 10^{-4}$ . We used a cross-entropy loss with label smoothing by a factor of 0.1. Moreover, a dropout of 0.1 and a gradient clipping to 20 were applied. Furthermore, we generated the outputs with a beam search of size 5. We loaded 8 sentences per update, with a frequency of 64, which supposes an effective batch size of 512. Those audios longer than 14 seconds and sentences with more than 300 tokens weren't used during training.

## 6.3. Results

This section describes the results of evaluating the ST system on WinoST and its performance in terms of gender. We are also interested in evaluating ASR English transcriptions and perceive if they contain any gender bias.

**General ASR and ST Evaluation:** We use the standard WER and BLEU measures to report the ASR and ST performance, respectively in Table 5. Our results concur with the results in (Di Gangi et al., 2019a).

**Gender Bias Evaluation in ST:** Our main objective is evaluating the accuracy of the systems for each of the language pairs. The high accuracy demonstrates that the system is able to translate the gender of the entities correctly. We also report  $\Delta G$  and  $\Delta S$  in Table 4. Ideally, these values should be close to 0. High  $\Delta G$  indicates that the system translates males better, and high  $\Delta S$  denotes that the system tends to translate pro-stereotypical entities better than anti-stereotypical entities.

The English-to-German (en-de) system has the highest accuracy 51%. This system also shows the minor difference in treating males and females translations (lowest  $\Delta G$ , 1.7) and the minor difference in the pro-stereotypical and the anti-stereotypical entities (lowest  $\Delta S$ , 1.5). The surprising behaviour comes with the English-to-Italian (en-it) system, which has the lowest accuracy of 37.3%, but still performs reasonably towards the anti-stereotypical entities translations, with the second lowest  $\Delta S$  difference (5.6). However, the system still favors the male translations with a high  $\Delta G$  difference (23.6). Both English-to-Spanish (en-es) and English-to-French (en-fr) have similar accuracies (45.2 and 43.2, respectively). However, there is a big difference in the  $\Delta S$ , which is much higher in the case of en-es (25.7), showing higher bias towards male translations. With these accuracy results, we are showing that the four translation directions present a significant amount of bias and they are far from approaching gender parity in performance. Moreover, after manually investigating the translation outputs, we observe that some professions are not correctly translated. 'nurse' is always translated to the female version in en-es and en-it, and similarly, 'developer' is always translated to the male version in en-it and en-fr, showing that stereotypes are perpetuated in ST. Many inflection errors, as illustrated in Table 3, can occur due to these stereotyped translations. Example 1 shows an anti-stereotypical co-reference case of 'nurse'. One of the common errors happens when translating the gendered adjective or pronoun correctly according to the context while referencing the wrong gendered stereotyped profession, shown in example 2 and 4. Another common problem is mismatched pronouns, the translation of the noun contradicts the profession's translation, due to biased translation in one of them, e.g. example 3. Example 5 shows a biased translation with a neutral pronoun.

**Gender Bias in ST vs MT:** Using the S-Transformer, the gender accuracy in the four languages is lower than the reported accuracy of MT commercial systems in the original WinoMT paper (Stanovsky et al., 2019). The best reported accuracies from commercial systems reached 74.1% in en-de, 59.4 % in en-es, 63.6% in en-fr and 42.4% in en-it, while in ST case, it is lower for all language pairs as shown in Table 5.

This may be due to the fact that ST is much more challenging than MT, and lower system performance implies higher biases. This big gap is reduced when comparing in terms of  $\Delta G$  and  $\Delta S$ . In this case, ST becomes closer to MT (when comparing in absolute terms), showing even better results in:  $\Delta G$  for en-it (in MT, 27.8);  $\Delta S$  for en-de (in MT, 12.5) and en-it (in MT, 9.4).

**Gender Bias Evaluation in ASR:** ASR systems have demonstrated gender biases for female speakers outputs (Tatman, 2017). However, gender bias associated with the context has not been studied in ASR yet,

<sup>4</sup><https://github.com/neulab/xnmt>

<sup>5</sup><https://github.com/moses-smt/mosesdecoder>

Table 3: Examples of outputs for different sentences from the WinoST corpus. Words in **blue**, **red**, and **orange** indicate male, female and neutral entities, respectively.

	Source transcription	[Target lang.] Predicted translation	Phenomenon
1	<b>The nurse</b> sent the carpenter to the hospital because <b>he</b> was extremely cautious	[ES] <b>El enfermero</b> envió el carpintero al hospital porque era extremadamente consciente	The gender of “nurse” is correctly inferred from the coreference.
2	<b>The construction worker</b> asked the nurse for an examination because <b>she</b> was injured	[ES] <b>El trabajador de la construcción</b> le pidió a la enfermera una examinación porque estaba <b>herida</b>	“Construction worker” is stereotyped to its male inflection, with a mismatched adjective gender.
3	The construction worker asked <b>the nurse</b> for an examination because <b>he</b> was the best	[ES] El trabajador de la construcción le pidió a <b>la enfermera</b> una examinación porque <b>él</b> era <b>el mejor</b>	“Nurse” is stereotyped to its female inflection, with a mismatched pronoun and adjective gender.
4	<b>The farmer</b> asked the designer what <b>she</b> could do to help	[FR] <b>L’agriculteur</b> a demandé au designer ce qu’ <b>elle</b> pouvait faire pour aider	Although the pronoun is translated correctly, “farmer” has a biased translation form.
5	The writer wrote a book about <b>the carpenter</b> because <b>her</b> story is very moving	[FR] L’écrivain a écrit un livre sur <b>le charpentier</b> , parce que <b>son</b> histoire est très émouvante	Biased form for “carpenter” with a neutral possessive gender

Table 4: WinoMT Gender Evaluation for four language pairs. Acc.(% of instances the translation had the correct gender)(the higher the better)  $\Delta G$  notes difference in F1 score between masculine and feminine sentences (the higher the worse) and  $\Delta S$  notes difference in accuracy between pro/anti stereotypical sentences (the higher the worse).

ST	Acc. (↑)	$\Delta G$ (↓)	$\Delta S$ (↓)
<i>en-de</i>	51.0	1.7	1.5
<i>en-es</i>	45.2	25.7	12.3
<i>en-fr</i>	43.2	13.7	14.5
<i>en-it</i>	37.3	23.6	5.6

Table 5: WER and BLEU (%) scores for the MuST-C corpus.

Language	ASR (WER ↓)	ST (BLEU ↑)
<i>en-de</i>	24.24	17.8
<i>en-es</i>	24.76	21.9
<i>en-fr</i>	23.98	28.2
<i>en-it</i>	24.18	18.3

and WinoST allows this analysis. We may expect that ASR is less prone to show gender bias in contextual patterns because of the nature of the task, which inherently combines the purpose of acoustic and language modeling. The acoustic part does not consider long context information, but it tends to benefit from local context information (Sperber et al., 2018). However, the language modeling part takes into consideration the

long-range context, and thus it may induce bias (Bordia and Bowman, 2019; Basta and Costa-jussà, 2021).

For further analysis of employing WinoST for ASR gender bias evaluation, it is required to distinguish between the gender’s errors in transcriptions and the general ones. Therefore, we computed the global accuracy in WinoST for the ASR best system in Table 5, en-fr, and got a 74.5% accuracy. However, this global accuracy includes 680 misspelled professions. Discarding these misspelling errors, we obtained a 98.72% accuracy of predicting pronouns, showing that the amount of gender bias at the context level is quite low in ASR.

## 7. Conclusions

This paper presents a new freely available challenge set for evaluating gender bias in ST. This challenge set, WinoST, can benefit from the evaluation protocol, which is widely used for MT. Our set is only based on evaluating the gender inaccuracies in translations in ST systems, mainly relying on the gender information extracted from the context and not from the audio signal. We used an S-Transformer end-to-end ST system and evaluated their accuracy in terms of gender bias with this new challenge set. Results show that gender accuracy is much lower for ST than for MT, but we have to take into account that ST has also a lower quality than MT. Finally, we show that ASR can exhibit a little gender bias at the contextual-level.

WinoST shares similar limitations as WinoMT, which is the fact of using a synthetic challenge set. Having a synthetic set is positive because of providing a controlled evaluation, and also it is negative because

we might be introducing some artificial biases. Therefore, further work could find in the wild transcriptions (with parallel speech utterances) that hold the valuable patterns designed in WinoMT, following (Levy et al., 2021).

As mentioned, future work includes making a challenge set from real-life examples, which excerpts human data bias. These examples should be covered by different speakers for studying the effect of gender of such speakers. Following the study of correlation of BLEU and gender bias accuracy (Kocmi et al., 2020), it would be useful to study the correlation between ASR quality and gender bias accuracy. Another approach is to study whether the accent of the speaker imposes bias in the translations. We assume that different accents may have an impact on the ST system. In this paper, we have focused on creating a benchmark for gender bias assessment, we are considering binary gender only. However, there is a need in the community for creating benchmarks for other types of biases, e.g. non-binary gender, racial bias (Koencke et al., 2020).

## 8. Acknowledgements

This work is supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 947657), the Catalan Agency for Management of University and Research Grants (AGAUR) through the FI PhD Scholarship, Universitat Politècnica de Catalunya with the collaboration of Banco de Santander and the Spanish Ministerio de Ciencia e Innovación through the project PID2019-107579RB-I00 / AEI / 10.13039/501100011033.

## 9. References

- Aharoni, R., Johnson, M., and Firat, O. (2019). Massively multilingual nmt. In *NAACL*, pages 3874–3884.
- Basta, C. and Costa-jussà, M. R. (2021). Impact of gender debiased word embeddings in language modeling. *CoRR*, abs/2105.00908.
- Basta, C., Costa-jussà, M. R., and Fonollosa, J. A. R. (2020). Towards mitigating gender bias in a decoder-based neural machine translation model by adding contextual information. In *4th WiNLP*, pages 99–102.
- Bentivogli, L., Savoldi, B., Negri, M., Di Gangi, M. A., Cattoni, R., and Turchi, M. (2020). Gender in danger? evaluating speech translation technology on the MuST-SHE corpus. In *ACL*, pages 6923–6933.
- Bérard, A., Pietquin, O., Besacier, L., and Servan, C. (2016). Listen and Translate: A Proof of Concept for End-to-End Speech-to-Text Translation. In *NIPS Workshop on end-to-end learning for speech and audio processing*.
- Bérard, A., Besacier, L., Kocabiyikoglu, A. C., and Pietquin, O. (2018). End-to-End ASR of Audio-books. In *ICASSP*.
- Blodgett, S. L., Barocas, S., Daumé III, H., and Wallach, H. (2020). Language (technology) is power: A critical survey of “bias” in NLP. In *ACL*, pages 5454–5476.
- Blodgett, S. L., Lopez, G., Olteanu, A., Sim, R., and Wallach, H. (2021). 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)*, pages 1004–1015, Online, August. Association for Computational Linguistics.
- Bordia, S. and Bowman, S. R. (2019). Identifying and reducing gender bias in word-level language models. In *NAACL: SRW*, pages 7–15.
- Cho, W. I., Kim, J. W., Kim, S. M., and Kim, N. S. (2019). On measuring gender bias in translation of gender-neutral pronouns. In *GeBNLP Workshop*, pages 173–181.
- Costa-jussà, M. R. and de Jorge, A. (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), December. Association for Computational Linguistics.
- Costa-jussà, M. R., Lin, P. L., and España-Bonet, C. (2020). Gebiotookit: Automatic extraction of gender-balanced multilingual corpus of wikipedia biographies. In *LREC*.
- Costa-jussà, M. R. (2019). An analysis of gender bias studies in natural language processing. *Nature Machine Intelligence*, 1.
- Di Gangi, M. A., Cattoni, R., Bentivogli, L., Negri, M., and Turchi, M. (2019a). Must-c: a multilingual speech translation corpus. In *NAACL*.
- Di Gangi, M. A., Negri, M., and Turchi, M. (2019b). Adapting Transformer to End-to-End Spoken Language Translation. In *Interspeech*, pages 1133–1137.
- D’Ignazio, C. and Klein, L. (2018). *Data Feminism*. MIT Press.
- Dong, L., Xu, S., and Xu, B. (2018). Speech-transformer: A no-recurrence sequence-to-sequence model for speech recognition. In *ICASSP*, pages 5884–5888.
- Font, J. E. and Costa-jussà, M. R. (2019). Equalizing gender bias in nmt with word embeddings techniques. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 147–154, Florence, Italy.
- Gonen, H. and Webster, K. (2020). Automatically identifying gender issues in machine translation using perturbations. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1991–1995.
- Kocmi, T., Limisiewicz, T., and Stanovsky, G. (2020).

- Gender coreference and bias evaluation at wmt 2020. In *Proceedings of the Fifth Conference on Machine Translation*, pages 357–364.
- Koenecke, A., Nam, A., Lake, E., Nudell, J., Quartey, M., Mengesha, Z., Touns, C., Rickford, J. R., Jurafsky, D., and Goel, S. (2020). Racial disparities in automated speech recognition. *Proceedings of the National Academy of Sciences*, 117(14):7684–7689.
- Levy, S., Lazar, K., et al. (2021). Collecting a large-scale gender bias dataset for coreference resolution and machine translation. In *EMNLP findings, 2021*.
- Liu, Y., Xiong, H., Zhang, J., and et al. (2019). End-to-End Speech Translation with Knowledge Distillation. In *Interspeech*, pages 1128–1132.
- Moryossef, A., Aharoni, R., and Goldberg, Y. (2019). Filling gender & number gaps in neural machine translation with black-box context injection. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 49–54.
- Neubig, G., Sperber, M., Wang, X., Felix, M., Matthews, A., Padmanabhan, S., Qi, Y., Sachan, D., Arthur, P., Godard, P., Hewitt, J., Riad, R., and Wang, L. (2018). XNMT: The eXtensible neural machine translation toolkit. In *AMTA*, pages 185–192.
- Prates, M. O., Avelar, P. H., and Lamb, L. C. (2020). Assessing gender bias in mt: a case study with google translate. *Neural Comput & Applic.*
- Saunders, D. and Byrne, B. (2020a). 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, ACL 2020, Online, July 5-10, 2020*, pages 7724–7736. Association for Computational Linguistics.
- Saunders, D. and Byrne, B. (2020b). Reducing gender bias in nmt as a domain adaptation problem. In *ACL*, pages 7724–7736.
- Sennrich, R., Haddow, B., and Birch, A. (2016). Nmt of rare words with subword units. In *ACL*, pages 1715–1725.
- Sperber, M., Niehues, J., Neubig, G., Stüker, S., and Waibel, A. (2018). Self-Attentional Acoustic Models. In *InterSpeech*.
- Stanovsky, G., Smith, N. A., and Zettlemoyer, L. (2019). Evaluating gender bias in mt. In *ACL*.
- Tatman, R. (2017). Gender and dialect bias in YouTube’s automatic captions. In *Proc. 1st ACL Workshop on Ethics in NLP*, pages 53–59.
- Vanmassenhove, E., Hardmeier, C., and Way, A. (2018). Getting gender right in nmt. In *EMNLP*.
- Vila, L., Escolano, C., Fonollosa, J. A. R., and Costajussà, M. R. (2018). End-to-End Speech Translation with the Transformer. In *IberSPEECH*, pages 60–63.
- Weiss, R. J., Chorowski, J., Jaitly, N., Wu, Y., and Chen, Z. (2017). Sequence-to-sequence models can directly translate foreign speech. In *Interspeech*, pages 2625–2629.