

# Breeding Gender-aware Direct Speech Translation Systems

Marco Gaido<sup>1,2†</sup>, Beatrice Savoldi<sup>2†</sup>, Luisa Bentivogli<sup>1</sup>, Matteo Negri<sup>1</sup>, Marco Turchi<sup>1</sup>

<sup>1</sup>Fondazione Bruno Kessler, Trento, Italy

<sup>2</sup>University of Trento, Italy

{mgaido, bentivo, negri, turchi}@fbk.eu, beatrice.savoldi@unitn.it

## Abstract

In automatic speech translation (ST), traditional *cascade* approaches involving separate transcription and translation steps are giving ground to increasingly competitive and more robust *direct* solutions. In particular, by translating speech audio data without intermediate transcription, direct ST models are able to leverage and preserve essential information present in the input (e.g. speaker’s vocal characteristics) that is otherwise lost in the cascade framework. Although such ability proved to be useful for gender translation, direct ST is nonetheless affected by gender bias just like its cascade counterpart, as well as machine translation and numerous other natural language processing applications. Moreover, direct ST systems that exclusively rely on vocal biometric features as a gender cue can be unsuitable and potentially harmful for certain users. Going beyond speech signals, in this paper we compare different approaches to inform direct ST models about the speaker’s gender and test their ability to handle gender translation from English into Italian and French. To this aim, we manually annotated large datasets with speakers’ gender information and used them for experiments reflecting different possible real-world scenarios. Our results show that gender-aware direct ST solutions can significantly outperform strong – but gender-unaware – direct ST models. In particular, the translation of gender-marked words can increase up to 30 points in accuracy while preserving overall translation quality.

## 1 Introduction

Language use is intrinsically social and situated as it varies across groups and even individuals (Bamman et al., 2014). As a result, the language data that are collected to build the corpora on which natural language processing models are trained are often far from being homogeneous and rarely offer a fair representation of different demographic groups and their linguistic behaviours (Bender and Friedman, 2018). Consequently, as predictive models learn from the data distribution they have seen, they tend to favor the demographic group most represented in their training data (Hovy and Spruit, 2016; Shah et al., 2020).

This brings serious social consequences as well, since the people who are more likely to be underrepresented within datasets are those whose representation is often less accounted for within our society. A case in point regards the gender data gap.<sup>1</sup> In fact, studies on speech taggers (Hovy and Søgaard, 2015) and speech recognition (Tatman, 2017) showed that the underrepresentation of female speakers in the training data leads to significantly lower accuracy in modeling that demographic group.

The problem of gender-related differences has also been inspected within automatic translation, both from text (Vanmassenhove et al., 2018) and from audio (Bentivogli et al., 2020). These studies – focused on the translation of spoken language – revealed a systemic gender bias whenever systems are required to overtly and formally express speaker’s gender in the target languages, while translating from languages that do not convey such information. Indeed, languages with grammatical gender, such as French and

---

<sup>†</sup>The authors contributed equally.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

<sup>1</sup>For a comprehensive overview on such societal issue see (Criado-Perez, 2019).

Italian, display a complex morphosyntactic and semantic system of gender agreement (Hockett, 1958; Corbett, 1991), relying on feminine/masculine markings that reflect speakers' gender on numerous parts of speech whenever they are talking about themselves (e.g. En: *I've never **been** there* – It: *Non ci sono mai **stata/stato***). Differently, English is a natural gender language (Hellinger and Bußman, 2001) that mostly conveys gender via its pronoun system, but only for third-person pronouns (*he/she*), thus to refer to an entity other than the speaker. As the example shows, in absence of contextual information (e.g. *As a woman, I have never been there*) correctly translating gender can be prohibitive. This is the case of traditional text-to-text machine translation (MT) and of the so-called *cascade* approaches to speech-to-text translation (ST), which involve separate transcription and translation steps (Stentiford and Steer, 1988; Waibel et al., 1991). Instead, *direct* approaches (Bérard et al., 2016; Weiss et al., 2017) translate without intermediate transcriptions. Although this makes them partially capable of extracting useful information from the input (e.g. by inferring speaker's gender from his/her vocal characteristics), the general problem persists: since female speakers (and associated feminine marked words) are less frequent within the training corpora, automatic translation tends towards a masculine default.

Following (Crawford, 2017), this attested systemic bias can directly affect the users of such technology by diminishing their gender identity or further exacerbating existing social inequalities and access to opportunities for women. Systematic gender representation problems – although unintended – can affect users' self-esteem (Bourguignon et al., 2015), especially when the linguistic bias is shaped as a perpetuation of stereotypical gender roles and associations (Levesque, 2011). Additionally, as the system does not perform equally well across gender groups, such tools may not be suitable for women, excluding them from benefiting from new technological resources.

To date, few attempts have been made towards developing gender-aware translation models, and surprisingly, almost exclusively within the MT community (Vanmassenhove et al., 2018; Elaraby et al., 2018; Moryossef et al., 2019). The only work on gender bias in ST (Bentivogli et al., 2020) proved that direct ST has an advantage when it comes to speaker-dependent gender translation (as in *I've never **been** there* uttered by a woman), since it can leverage acoustic properties from the audio input (e.g. speaker's fundamental frequency). However, relying on perceptual markers of speakers' gender is not the best solution for all kinds of users (e.g. transgenders, children, vocally-impaired people). Moreover, although their conclusions remark that direct ST is nonetheless affected by gender bias, no attempt has yet been made to try and enhance its gender translation capability. Following these observations, and considering that ST applications have entered widespread societal use, we believe that more effort should be put into further investigating and controlling gender translation in direct ST, in particular when the gender of the speaker is known in advance.

Towards this objective, we annotated MuST-C (Di Gangi et al., 2019a; Cattoni et al., 2020) - the largest freely available multilingual corpus for ST - with speakers' gender information and explored different techniques to exploit such information in direct ST. The proposed techniques are compared, both in terms of overall translation quality as well as accuracy in the translation of gender-marked words, against a "pure" model that solely relies on the speakers' vocal characteristics for gender disambiguation.

In light of the above, our contributions are:

- (1) the manual annotation of the TED talks contained in MuST-C with speakers' gender information, based on the personal pronouns found in their TED profile. The resource is released under a CC BY NC ND 4.0 International license, and is freely downloadable at <https://ict.fbk.eu/must-speakers/>;
- (2) the first comprehensive exploration of different approaches to mitigate gender bias in direct ST, depending on the potential users, the available resources and the architectural implications of each choice.

Experiments carried out on English-Italian and English-French show that, on both language directions, our gender-aware systems significantly outperform "pure" ST models in the translation of gender-marked words (up to 30 points in accuracy) while preserving overall translation quality. Moreover, our best systems learn to produce feminine/masculine gender forms regardless of the perceptual features received from the audio signal, offering a solution for cases where relying on speakers' vocal characteristics is detrimental to a proper gender translation.

## 2 Background

Besides the abundant work carried out for English monolingual NLP tasks (Sun et al., 2019), a consistent amount of studies have now inspected how MT is affected by the problem of gender bias. Most of them, however, do not focus on speaker-dependent gender agreement. Rather, a number of studies (Stanovsky et al., 2019; Escudé Font and Costa-jussà, 2019; Saunders and Byrne, 2020) evaluate whether MT is able to associate pronominal coreference with an occupational noun to produce the correct masculine/feminine forms in the target gender-inflected languages (En: *I've known her for a long time, **my friend** is a cook.* Es: *La conozco desde hace mucho tiempo, **mi amiga** es cocinera).*

Notably, few approaches have been employed to make neural MT systems speaker-aware by controlling gender realization in their output. Elaraby et al. (2018) enrich their data with a set of gender-agreement rules so to force the system to account for them in the prediction step. In (Vanmassenhove et al., 2018), the MT system is augmented at training time by prepending a gender token (female or male) to each source segment. Similarly, Moryossef et al. (2019) artificially inject a short phrase (e.g. *she said*) at inference time, which acts as a gender domain label for the entire sentence. These approaches are implemented and tested on natural spoken language that, compared to written language, is more likely to contain references to the speaker and, consequently, speaker-dependent gender-marked words.

In the light of above, the correct translation of gender is a particularly relevant task for ST systems, as they are precisely developed to translate oral, conversational language. Nonetheless, to our knowledge only one work has investigated gender bias in ST (Bentivogli et al., 2020). Focusing on the proper handling of gender phenomena, the authors take stock of the situation by comparing cascade and direct architectures on MuST-SHE, a multilingual benchmark derived from the TED-based MuST-C corpus and specifically designed to evaluate gender translation and bias in ST. Their conclusions remark that, although traditional cascade systems still outperform direct solutions, the latter are able to exploit audio information for a better treatment of speaker-dependent gender phenomena.

These findings open a line of focused research on speaker-aware ST that is worth exploring more thoroughly, also in light of the fact that the performance gap between cascade and direct approaches has further reduced (Ansari et al., 2020). On one side, rather than comparing the two paradigms, this progress now motivates exploring all the possible ways to boost direct ST performance towards the translation of gender-marked expressions. On the other side, since the direct systems tested in (Bentivogli et al., 2020) rely on “pure” models built to verify an hypothesis (i.e. that translating audio signals without intermediate representations makes a difference in handling gender), the real potential of direct ST technology with respect to this problem is still unknown. Moreover, as their “pure” models solely rely on the speaker’s fundamental frequency, various instances in which such perceptual marker is not indicative of the speaker’s gender remain out of the picture.

## 3 Annotation of MuST-C with Speakers’ Gender Information

Although current research on gender-aware ST can count on the MuST-SHE benchmark (Bentivogli et al., 2020) for fine-grained evaluations, gender-annotated training data are not yet available. So far, this has limited the scope of research to application scenarios in which speakers’ gender is inferred from the input audio. These scenarios are not representative of the full range of possible usages of ST and are also potentially problematic, since gendered forms expected in translation do not necessarily align with speaker’s vocal characteristics.

In the light of the above, building large training corpora explicitly annotated with gender information becomes crucial. To this aim, rather than building a new resource from scratch, we opted for adding an annotation layer to MuST-C, which has been chosen over other existing corpora (Iranzo-Sánchez et al., 2020) for the following reasons: *i*) it is currently the largest freely available multilingual corpus for ST, *ii*) being based on TED talks it is the most compatible one with MuST-SHE, *iii*) TED speakers’ personal information is publicly available and retrievable on the TED official website.<sup>2</sup>

Following the MuST-C talk IDs, we have been able to *i*) automatically retrieve the speakers’ name,

<sup>2</sup>Available at <https://www.ted.com/speakers/>

	Talks M	Talks F	Hours M	Hours F	Segments M	Segments F
<b>en-it</b>	1,569	725	316	136	178,841	71,877
<b>en-fr</b>	1,569	725	327	151	189,742	81,527

Table 1: Statistics for MuST-C data with gender annotation. The number of segments and hours varies over the two language pairs due to the different pre-processing of MuST-C data.

*ii*) find their associated TED official page, and *iii*) manually label the personal pronouns used in their descriptions. Though time-consuming, such manual retrieval of information is preferable to automatic speaker gender identification for the following reasons. First, since automatic methods based on fundamental frequency are not equally accurate across demographic groups (e.g. women and children are hard to distinguish as their pitch is typically high (Levitan et al., 2016)), manual assignment prevents from incorporating gender misclassifications in our training data. Second, biological essentialist frameworks that categorize gender based on acoustic cues (Zimman, 2020) are especially problematic for transgender individuals, whose gender identity is not aligned with the sex they have been assigned at birth based on designated anatomical/biological criteria (Stryker, 2008).

Differently, following the guidelines in (Larson, 2017), we do not want to run the risk of making assumptions about speakers’ gender identity and introducing additional bias within an environment that has been specifically designed to inspect gender bias. By looking at the personal pronouns used by the speakers to describe themselves, our manual assignment instead is meant to account for the gender linguistic forms by which the speakers accept to be referred to in English (GLAAD, 2007), and would want their translations to conform to. We stress that gendered linguistic expressions do not directly map to speakers’ self-determined gender identity (Cao and Daumé III, 2020). We therefore make explicit that throughout the paper, when talking about speakers’ gender, we refer to their accepted linguistic expression of gender rather than their gender identity.

Focusing on the two language pairs of our interest, 2,294 different speakers described via *he/she* pronouns<sup>3</sup> are represented in both en-it and en-fr. Their male/female<sup>4</sup> distribution is unbalanced, as shown in Table 1, which presents the number of talks, as well as the total number of segments and the corresponding hours of speech.

## 4 ST Systems

For our experiments, we built three types of direct systems. One is the *base* system, a state-of-the-art model that does not leverage any external information about speaker’s gender (§4.1). The others are two gender-aware systems that exploit speakers’ gender information in different ways: *multi-gender* (§4.2) and *specialized* (§4.3). All the models share the same architecture, a Transformer (Vaswani et al., 2017) adapted to ST. The encoder processes the input Mel-filter-bank sequences with two 2D convolutional layers with stride 2, returning a sequence that is four times shorter than the original input. The vectors of this sequence are projected by a linear transformation into the dimensional space used in the following encoder Transformer layers and are summed with sinusoidal positional embeddings. The attentions in the encoder layers are biased toward elements close on the time dimension with a logarithmic distance penalty (Di Gangi et al., 2019b). The decoder architecture, instead, is not modified.

<sup>3</sup>It is important to point out that some individuals do not neatly fall into the female/male binary (gender fluid, non-binary) or may even not experience gender at all (a-gender) (Richards et al., 2016; Schilt and Westbrook, 2009; GLAAD, 2007), possibly preferring the use of singular *they* or other neopronouns. Within MuST-C, speakers with *they* pronoun have been encountered, but MuST-C human-reference translations do not exhibit linguistic gender-neutralization strategies, which are difficult to fully implement in languages with grammatical gender (Lessinger, 2020). Note that, because of such inconsistency and the very limited number of cases, these instances were not used for training. Our experiments therefore focus on binary linguistic forms. By design, some sparse talks with multiple speakers of different genders were also excluded. Detailed information about all MuST-C speakers and corresponding talks can be found in the resource release at [ict.fbk.eu/must-speakers](http://ict.fbk.eu/must-speakers).

<sup>4</sup>Some authors distinguish female/male for sex and woman/man for gender (among others (Larson, 2017)). For the sake of simplicity, in our study we use female/male to respectively indicate those speakers whose personal pronouns are *she/he*.

## 4.1 Base ST Model

We are interested in evaluating and improving gender translation on strong ST models that can be used in real-world contexts. As such, our base, gender-unaware model is trained with the goal of achieving state-of-the-art performance on the ST task. To this aim, we rely on data augmentation and knowledge transfer techniques that were shown to yield competitive models at the IWSLT-2020 evaluation campaign (Ansari et al., 2020; Potapczyk and Przybysz, 2020; Gaido et al., 2020). In particular, we use three data augmentation methods – SpecAugment (Park et al., 2019), time stretch (Nguyen et al., 2020), and synthetic data generation (Jia et al., 2019) – and we transfer knowledge both from ASR and MT through component initialization and knowledge distillation (Hinton et al., 2015).

The ST model’s encoder is initialized with the encoder of an English ASR model (Bansal et al., 2019) with a lower number of encoder layers (the missing layers are initialized randomly, as well as the decoder). This ASR model is trained on Librispeech (Panayotov et al., 2015), Mozilla Common Voice,<sup>5</sup> How2 (Sanabria et al., 2018), TEDLIUM-v3 (Hernandez et al., 2018), and the utterance-transcript pairs of the ST corpora – EuroParl-ST (Iranzo-Sánchez et al., 2020) and MuST-C. These datasets are either gender unbalanced or do not provide speaker’s gender information apart from Librispeech, which is balanced in terms of female/male speakers (Garnerin et al., 2020). However, since these speakers are just book narrators, first-person sentences do not really refer to the speakers themselves.

Knowledge distillation (KD) is performed from a *teacher* MT model by optimizing the cross entropy between the distribution produced by the teacher and by the *student* ST model being trained (Liu et al., 2019). For both en-it and en-fr, the MT model is trained on the OPUS datasets (Tiedemann, 2016).

The ST model is trained in three consecutive steps. In the first step, we use the synthetic data obtained by pairing ASR audio samples with the automatic translations of the corresponding transcripts. In the second step, the model is trained on the ST corpora. In these first two steps, we use the KD loss function. Finally, in the third step, the model is fine-tuned on the same ST corpora using label-smoothed cross entropy (Szegedy et al., 2016). SpecAugment and time stretch are used in all steps.

## 4.2 Multi-gender Systems

The idea of “multi-gender” models, i.e. models informed about the speaker’s gender with a tag prepended to the source sentence, was introduced by Vanmassenhove et al. (2018) and Elaraby et al. (2018). This approach was inspired by one-to-many multilingual neural MT systems (Johnson et al., 2017), in which a single model is trained to translate from a source into many target languages by means of a *target-forcing* mechanism. With this mechanism - here adapted for “*gender-forcing*” - ST multi-gender systems are fed not only with the input audio, but also with a tag (*token*) representing the speaker’s gender. This *token* is converted into a vector through learnable embeddings. This approach has two main potential advantages: *i*) a single model supports both male and female speakers (which makes it particularly appealing for real-world application scenarios), and *ii*) each gender direction can benefit from the data available for the other, potentially learning to produce words that would have never been seen otherwise (*transfer learning*). Regarding the several options to supply the model with the additional gender information, we do not follow the approach of Vanmassenhove et al. (2018) and Elaraby et al. (2018), since it is dedicated to MT. Instead, we consider those that obtained the best results in multilingual direct ST (Di Gangi et al., 2019c; Inaguma et al., 2019), namely:

**Decoder prepending.** The gender token replaces the  $\langle \backslash s \rangle$  (*EOS*, end-of-sentence) that is added in front of the generated tokens in the decoder input.

**Decoder merge.** The gender embedding is added to all the word embeddings representing the generated tokens in the decoder input.

**Encoder merge.** The gender embedding is added to the Mel-filter-bank sequence representing the source speech given as input to the encoder.

In all cases, multi-gender models’ weights are initialized with those of the *Base* models. The only randomly-initialized parameters are those of the gender embeddings.

---

<sup>5</sup><https://voice.mozilla.org/>

### 4.3 Gender-specialized Systems

In this approach, two different gender-specific models are created. Each model is initialized with the *Base* model’s weights and then fine-tuned only on samples of the corresponding speaker’s gender. This solution has the drawback of a higher maintenance burden than the multi-gender one, as it requires the training and management of two separate models. Moreover, no transfer learning is possible: although each model is initialized with the base model trained on all the data and the low learning rate used in the fine-tuning prevents catastrophic forgetting (McCloskey and Cohen, 1989), data scarcity conditions for a specific gender are likely to lead to lower performance on that direction.

### 4.4 Gender-balanced Validation Set

To train our gender-aware models, we do not rely on the standard MuST-C validation set as it reflects the same gender-imbalanced distribution found in the training data. We therefore created a new specifically designed validation set composed of 20 talks. Unlike the standard MuST-C validation set, it contains a balanced number of female/male speakers, thus avoiding to reward models’ potentially biased behaviour. This new resource is released under a CC BY NC ND 4.0 International license, and is freely downloadable at <https://ict.fbk.eu/must-c-gender-dev-set/>.<sup>6</sup>

## 5 Experimental Setting

### 5.1 Experiments

As described in §4.1, our ST models adopt knowledge transfer techniques that showed to significantly improve ST performance. In particular, knowledge distillation (KD) is especially relevant as it allows the ST model to learn and exploit the wealth of training data available for MT, which otherwise would not be accessible. Hence, since we are also interested in assessing the effect of KD on the ability of the resulting ST systems to deal with gender, we compare: *i*) the teacher MT models, *ii*) the intermediate ST models trained on KD, and *iii*) the final ST models obtained with fine-tuning without KD.

The final ST models are used to initialize both multi-gender (§4.2) and gender-specialized models (§4.3), which are then fine-tuned on the MuST-C gender-labeled dataset. Since, as seen in §3, this dataset shows a quite skewed male/female speaker distribution (approximately 70%/30%), we test both approaches in two different data conditions: *i*) balanced (\*-BAL), where we use all the female data available together with a random subset of the male data, and *ii*) unbalanced (\*-ALL) where all the MuST-C data available are exploited. It must be noted that there are differences between the two approaches on the usage of data. In the specialized approach, since we have two separate systems, the one which is fine-tuned with talks by female speakers remains the same in both data conditions. Differently, in the multi-gender approach, which is trained on both genders together, all the training mini-batches contain the same number of samples for each gender. Thus, when all MuST-C data are used, the female gender pairs – which are underrepresented – are over-sampled.

### 5.2 Evaluation Method

For our experiments, we rely on MuST-SHE (Bentivogli et al., 2020), a gender-sensitive, multilingual benchmark for MT and ST consisting of (*audio, transcript, translation*) aligned triplets. By design, each segment in the corpus requires the translation of at least one English gender-neutral word into the corresponding masculine/feminine target word(s) to convey a referent’s gender. With the intent to evaluate our gender-aware ST models on speaker-dependent gender phenomena, we focus on a portion of MuST-SHE containing, for each language pair, ~600 segments where gender agreement only depends on the speaker’s gender.<sup>7</sup> Segments are balanced with respect to female/male speakers and masculine/feminine marked words, which are explicitly annotated in the corpus.

An important feature of MuST-SHE is that, for each reference translation, an almost identical “wrong” reference is created by swapping each annotated gender-marked word into its opposite gender (e.g. *I*

<sup>6</sup>To ease future research on gender bias in ST for the three language pairs represented in MuST-SHE (en-it, en-fr, en-es), the validation set is also available for en-es.

<sup>7</sup>In (Bentivogli et al., 2020) this portion is referred to as “Category 1”.

have been uttered by a woman is translated into the correct Italian reference *Sono stata*, and into the wrong reference *Sono stato*). The idea behind gender-swapping is that the difference between the scores computed against the “correct” and the “wrong” reference sets captures the system’s ability to handle gender translation. However, relying on these scores does not allow to distinguish between those cases where the system “fails” by producing a word different from the one present in the references (e.g. *andar\** in place of *stat\**) and failures specifically due to the wrong realization of gender (e.g. *stato* in place of *stata*).

Thus, while following the same principles as Bentivogli et al. (2020), in our experiments we rely on a more informative evaluation. First, we calculate the **term coverage** as the proportion of gender-marked words annotated in MuST-SHE that are actually generated by the system, on which the accuracy of gender realization is therefore *measurable*. Then, we define **gender accuracy** as the proportion of correct gender realizations among the words on which it is *measurable*. Our evaluation method has several advantages. On one side, *term coverage* unveils the precise amount of words on which systems’ gender realization is measurable. On the other, *gender accuracy* directly informs about systems’ performance on gender translation and related gender bias: scores below 50% indicate that the system produces the wrong gender more often than the correct one, thus signalling a particularly strong bias. Gender accuracy has the further advantage of informing about the margins for improvement of the systems.

## 6 Results

### 6.1 Overall Results

Table 2 presents overall results in terms of BLEU scores on the MuST-SHE test set. Despite the well-known differences in performance between en-it and en-fr, both language directions show the same trend.

First, the MT systems used by the ST models for KD achieve by far the highest performance. This is expected since the ST task is more complex and MT models are trained on larger amounts of data. However, all our ST results are competitive compared to those published for the two target languages. In particular, on the MuST-C test set, the scores of our ST BASE models are 27.7 (en-it) and 40.3 (en-fr), respectively 0.3 and 4.8 BLEU points above the best *cascade* results reported in (Bentivogli et al., 2020).

Moving on to ST systems, we attest that the models after the first two training steps based on KD (BASE-KD-ONLY, see 4.1) have a lower translation quality than the BASE models, showing that the third training step is crucial to boost overall performance. In general, except for the MULTI-DECMERGE system (whose performance is significantly lower), we do not observe statistically significant differences between the BASE models and their gender-aware extensions (MULTI-\* and SPECIALIZED-\*), which also perform on par when fine-tuned with varying amounts of annotated data (balanced vs all).

Due to the very small percentage of speaker-dependent gender-marked words in MuST-SHE (< 3%, 810-840 over ~30,000 words), systems’ ability to translate gender is not reflected by BLEU scores. Now, we delve deeper into our more informative evaluation (as per §5.2) and turn to the term coverage and gender accuracy values presented in Table 3. The overall results assessed with BLEU are confirmed by **term coverage** scores for both en-it and en-fr: the MT systems generate the highest number of annotated

	En-It	En-Fr
	Bleu	Bleu
MT for KD	33.59	39.61
Base-KD-only	23.58	31.97
Base	27.51	34.25
Multi-DecPrep-Bal	26.36	33.54
Multi-DecPrep-All	26.17	34.13
Multi-EncMerge-Bal	26.47	33.29
Multi-EncMerge-All	26.39	33.07
Multi-DecMerge-Bal	21.99	27.06
Multi-DecMerge-All	22.12	27.74
Specialized-Bal	27.43	34.32
Specialized-All	27.79	34.61

Table 2: BLEU scores on MuST-SHE.

	En-It		En-Fr	
	Cover.	Acc.	Cover.	Acc.
MT for KD	63.83	51.45	63.10	52.08
Base-KD-only	56.05	51.76	59.17	53.12
Base	56.17	56.26	62.02	56.24
Multi-DecPrep-Bal	56.91	64.86	60.95	69.34
Multi-DecPrep-All	56.54	66.81	61.31	70.29
Multi-EncMerge-Bal	57.04	62.55	60.60	62.67
Multi-EncMerge-All	57.65	60.39	62.38	61.83
Multi-DecMerge-Bal	49.88	59.41	54.52	64.63
Multi-DecMerge-All	50.74	60.58	56.31	65.96
Specialized-Bal	57.90	86.35	61.79	86.13
Specialized-All	58.02	87.02	62.38	86.45

Table 3: Term coverage and gender accuracy scores.

words present in MuST-SHE (63.83% on en-it and 63.10% on en-fr), while we do not observe large differences among the ST models (between 56.17% and 58.02% for en-it and 60.60% and 62.38% for en-fr). Instead, looking at **gender accuracy**, we immediately unveil that overall performance is not an indicator of the systems’ ability to translate gender. In fact, the best performing MT systems show the lowest gender accuracy (51.45% for en-it and 52.08% for en-fr): intrinsically constrained by the lack of access to audio information, they produce the wrong target gender in half of the cases. Such deficiency is directly reflected in the BASE-KD-ONLY models, which are strongly influenced by the MT behaviour; thus, although effective for overall quality, KD is detrimental to gender translation. By undergoing the third training step without KD, the BASE models are in fact able to improve on gender translation, but with limited gains. Differently, the models fed with the speaker’s gender information display a noticeable increase in gender translation, with SPECIALIZED-\* models outperforming the MULTI-\* ones by 16–20 points and the BASE ones by 30 points.

Among the multi-gender architectures, our results show that MULTI-DECPREP has an edge on the other two models, both in overall and gender translation performance: for the sake of simplicity, from now on we thus present only that model. As a single-model architecture, multi-gender would be a more functional solution than multiple specialized models, but – being trained on both female and male speakers’ utterances – it is noticeably weaker than multiple specialized models (trained on gender-specific data) at predicting gender. With regard to the different amounts of gender-annotated data used to train our gender-aware models, we cannot see any appreciable variation in term coverage and gender accuracy between the two settings. Further insights on this aspect are presented in the next section.

## 6.2 Cross-gender Analysis

Table 4 shows separate term coverage and gender accuracy scores for target feminine and masculine forms. This allows us to highlight the models’ translation ability for each gender form and conduct cross-gender comparisons to detect potential bias. Also in this analysis, results are consistent across language pairs. We assess that both the MT model and its strongly connected BASE-KD-ONLY present a very strong bias since they almost always produce masculine forms: accuracy is always much lower than 50% on the feminine set (up to 20.85% for en-it and 26.91% for en-fr) and very high on the masculine set (up to 88.49% for en-it and 89.58% for en-fr). After fine-tuning without KD, the BASE ST models improve feminine forms realization, but they still remain far from 50%. The comparison with the direct model in (Bentivogli et al., 2020) shows that, despite the much higher overall translation quality, our BASE models are affected by a stronger bias. This further confirms the detrimental effect of KD on gender translation and that higher overall quality does not directly imply a better speaker’s gender treatment.

All gender-aware models significantly reduce bias with respect to the BASE systems. This is particularly evident in the feminine set, where accuracy scores far above 50% indicate their ability to correctly represent female speakers. In particular, the SPECIALIZED models achieve the best results on both feminine and masculine sets (over 79% and 93% respectively). The higher performance on the masculine set can be explained considering that the two gender-specialized models derive from the BASE model, which is strongly biased towards masculine forms. Interestingly, MULTI-DECPREP shows similar feminine/masculine accuracy scores. This is possibly due to the random initialization of the gender tokens’ embeddings: as a result, the initial model hidden representations and predictions are perturbed in an

	En-it				En-Fr			
	Feminine		Masculine		Feminine		Masculine	
	Cover.	Acc.	Cover.	Acc.	Cover.	Acc.	Cover.	Acc.
MT for KD	66.25	16.23	61.46	88.49	63.76	16.24	62.41	89.58
Base-KD-only	58.75	20.85	53.41	84.93	58.59	26.91	59.76	79.44
Base	58.75	33.62	53.66	80.45	60.47	32.30	63.61	79.55
Multi-DecPrep-Bal	60.00	68.75	53.90	60.63	61.41	68.58	60.48	70.12
Multi-DecPrep-All	58.00	69.83	55.12	63.72	61.88	65.78	60.72	75.00
Specialized-Bal	62.00	79.84	53.90	93.67	62.59	79.32	60.96	93.28
Specialized-All	62.00	79.84	54.15	95.05	62.59	79.32	62.17	93.80

Table 4: Coverage and accuracy scores divided by feminine and masculine word forms.



unbiased way. An unbiased starting condition combined with balanced data leads to a fairer, similar behaviour across genders, although the final models have a lower accuracy than the SPECIALIZED ones.

Finally, we notice that results obtained by training our models with balanced (\*-BAL) and unbalanced (\*-ALL) datasets are similar. Indeed, the masculine gender accuracy slightly improves by adding more male data, while there is not a clear trend on the feminine accuracy: we can conclude that oversampling the data is functional inasmuch it keeps the performance on the feminine set stable.

### 6.3 Analysing Conflicts between Vocal Characteristics and Gender Tags

So far, we worked under the assumption that the speaker’s vocal characteristics match with those typically associated to the gender category she/he identifies with. In this section, we explore systems’ capacity to produce translations that are coherent with the speaker’s gender in a scenario in which this assumption does not hold: this is the case of some transgenders, children and people with vocal impairment. However, we are hindered by the almost absent representation of such users within MuST-C. As such, we design a counterfactual experiment where we associate the opposite gender tag to each actual female/male speaker and inspect models’ behaviour when receiving conflicting information between the gender tag and the properties of the acoustic signal. This can also be considered as an indirect assessment of systems’ robustness to possible errors in application scenarios where speakers’ gender is assigned automatically.

Table 5 presents the results for this experiment. In the *M-audio/F-transl* set, systems were fed with a male voice and a female tag and the expected translation is in the feminine form, while in the *F-audio/M-transl* set we have the opposite. As we can see, in both sets the multi-gender model has a drastic drop in accuracy with respect to the results shown in Table 4, with scores below 50% for en-it. This behaviour indicates that this model relies on both the gender token and the audio features, which in this scenario are conflicting. Thus, the multi-gender model could be more robust to possible errors in automatic recognition of the speaker’s gender, but it is not usable in scenarios in which the vocal characteristics have to be ignored. On the contrary, the specialized systems show a high accuracy on both sets. In particular, on *F-audio/M-transl* the performance is in line with the results of Table 4. This indicates that, independently from speakers’ vocal characteristics, the model relies only on the provided gender information, being therefore suitable for situations in which one wants to control the gendered forms in the output and override the potentially misleading speech signals.

	En-it				En-Fr			
	M-audio/F-transl.		F-audio/M-transl.		M-audio/F-transl.		F-audio/M-transl.	
	Cover.	Acc.	Cover.	Acc.	Cover.	Acc.	Cover.	Acc.
Multi-DecPrep-All	54.88	45.78	60.25	38.17	61.93	45.14	61.18	55.77
Specialized-All	54.39	64.57	60.75	94.24	62.17	59.69	61.41	94.25

Table 5: Coverage and accuracy scores when the correct translation is expected in a gender form opposite to the speaker’s gender but in accordance with the gender tag fed to the system.

## 7 Manual Analysis

We complement our automatic evaluation with a manual inspection on the output of three models: BASE, MULTI-DECPREP-ALL (MULTI), and SPECIALIZED-ALL (SPEC). For each model, we analyzed the translation of 100 common segments across en-it/en-fr, which allow for cross-lingual comparisons.

We first take into account those instances where systems’ accuracy in the production of gender-marked words was *measurable*, as in (a), (b), (c) in Table 6. A first observation, consistent across languages and models, is that a controlling noun (*student*) and its modifiers (*the, classic, Asian*) always concord in gender in the systems’ output. As per (a), this agreement is respected for both correct (MULTI, SPEC) and wrong gender realizations (BASE). Differently, (b) shows that, whenever two words are not related by any morphosyntactic dependency, some words may be correctly translated (*chercheuse* – MULTI, SPEC), and some others not (*professeur*). Such dynamic seems to attest that, although the systems are fed with sentence-level gender tags, gender predictions are still skewed at the level of the single word.

(a) F	IT	SRC	In one, I was <b>the classic Asian student</b> ...	
		REF	In uno ero <b>la classica studentessa asiatica</b> ...	
		BASE	In una, ero <b>il classico studente asiatico</b> ...	X X X X
		MULTI & SPEC	Innanzitutto, ero <b>la classica studentessa asiatica</b> ...	✓✓✓✓
(b) F	FR	SRC	As a <b>researcher</b> , a <b>professor</b> ...	
		REF	En tant que <b>chercheuse, professeure</b> ...	
		BASE	En tant que <b>chercheur, professeur</b> ...	X X
		MULTI & SPEC	En tant que <b>chercheuse, professeur</b> ...	✓X
(c) M	IT	SRC	...the <u>woman</u> who wanted to know me as an <b>adult</b> ...	
		REF	...la <u>donna</u> che voleva vedere come fossi diventato da <b>adulto</b> ...	
		BASE & MULTI	... una donna che voleva conoscermi da <b>adulto</b> ...	X
		SPEC	... una donna che voleva conoscermi da <b>adulto</b> ...	✓
(d) F	FR	SRC	When I was a <b>kid</b> ...	
		REF	Quand j'étais <b>petite</b> ...	
		BASE & MULTI & SPEC	Quand j' ai été <b>tuée</b> ...	?
(e) F	IT	SRC	...while downhill skiing <b>paralyzed</b> me.	
		REF	...quando un terrificante incidente sciistico mi ha lasciato <b>paralizzata</b> .	
		BASE & MULTI & SPEC	... quando mi <b>paralizzò</b> .	?
(f) F	IT	SRC	I was <b>elected</b> ...	
		REF	Sono stata <b>eletta</b> ...	
		BASE	Fui <b>eletto</b> ...	X
		MULTI & SPEC	Fui <b>selezionata</b> ...	?

Table 6: Examples of feminine (F) and masculine (M) gender-marked words translated by BASE, MULTI-DECPREP-ALL (MULTI) and SPECIALIZED-ALL (SPEC) on en-it and en-fr.

Overall, (a), (b) and (c) clearly attest the progressively improved performance from BASE to MULTI and SPEC. In particular, in (c), SPEC is able to pick the required masculine form in spite of a contextual hint about a second female referent (*woman*), thus overcoming what is a difficult prediction even for MULTI.

We also inspect those cases where systems' accuracy on gender production was not *measurable* to cast some light on the reasons for a limited *term coverage*. We found that, while there are some generally wrong translations – (d) – such instances only amount to 1/3 of the cases. In the remaining 2/3, the output is fluent and reflects the source utterance meaning but it simply does not match the exact annotated word in the reference. We found that ST translations often offer alternative constructions that do not require an overt gender-inflection – (e) – or rely on appropriate gender-marked synonyms of the word in the reference – (f). We can hence conclude that many gender translations that do not contribute to *gender accuracy* confirm an improved ability of the enriched models in gender translation.

## 8 Conclusion

We rose to the challenge posed by Bentivogli et al. (2020) to further explore gender translation in direct ST. Going beyond direct systems' attested ability to leverage speaker's vocal characteristics from the audio input, we developed gender-aware models suitable for operating conditions where speaker's gender is known. To this aim, we annotated the large MuST-C dataset with speaker's gender information, and used the new annotations to experiment with different architectural solutions: "multi-gender" and "specialized". Our results on two language pairs (en-it and en-fr) show that breeding speaker's gender-aware ST improves the correct realization of gender. In particular, our specialized systems outperform the gender-unaware ST models by 30 points in gender accuracy without affecting overall translation quality.

## Acknowledgements

We would like to thank the anonymous reviewers and the COLING'2020 Ethics Advisory Group for their insightful comments. This work is part of the "End-to-end Spoken Language Translation in Rich Data Conditions" project,<sup>8</sup> which is financially supported by an Amazon AWS ML Grant.

<sup>8</sup><https://ict.fbk.eu/units-hlt-mt-e2eslt/>

## References

- Ebrahim Ansari, Amittai Axelrod, Nguyen Bach, Ondřej Bojar, Roldano Cattoni, Fahim Dalvi, Nadir Durrani, Marcello Federico, Christian Federmann, Jiatao Gu, Fei Huang, Kevin Knight, Xutai Ma, Ajay Nagesh, Matteo Negri, Jan Niehues, Juan Pino, Elizabeth Salesky, Xing Shi, Sebastian Stüker, Marco Turchi, Alexander Waibel, and Changhan Wang. 2020. FINDINGS OF THE IWSLT 2020 EVALUATION CAMPAIGN. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 1–34, Online, July. Association for Computational Linguistics.
- David Bamman, Chris Dyer, and Noah A. Smith. 2014. Distributed Representations of Geographically Situated Language. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 828–834, Baltimore, Maryland, June. Association for Computational Linguistics.
- Sameer Bansal, Herman Kamper, Karen Livescu, Adam Lopez, and Sharon Goldwater. 2019. Pre-training on High-resource Speech Recognition Improves Low-resource Speech-to-text Translation. 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)*, pages 58–68, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Emily M. Bender and Batya Friedman. 2018. Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Luisa Bentivogli, Beatrice Savoldi, Matteo Negri, Mattia A. Di Gangi, Roldano Cattoni, and Marco Turchi. 2020. Gender in danger? Evaluating Speech Translation Technology on the MuST-SHE Corpus. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6923–6933, Online, July. Association for Computational Linguistics.
- David Bourguignon, Vincent Y. Yzerbyt, Catia P. Teixeira, and Ginette Herman. 2015. When does it hurt? Intergroup permeability moderates the link between discrimination and self-esteem. *European Journal of Social Psychology*, 45(1):3–9, 2.
- Alexandre Bérard, Olivier Pietquin, Christophe Servan, and Laurent Besacier. 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*, Barcelona, Spain, December.
- Yang T. Cao and Hal Daumé III. 2020. Toward Gender-Inclusive Coreference Resolution. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4568–4595, Online, July. Association for Computational Linguistics.
- Roldano Cattoni, Mattia A. Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2020. MuST-C: A Multilingual Corpus for end-to-end Speech Translation. *Computer Speech & Language Journal*. Doi: <https://doi.org/10.1016/j.csl.2020.101155>.
- Greville G. Corbett. 1991. *Gender*. Cambridge University Press, Cambridge, UK.
- Kate Crawford. 2017. The Trouble with Bias. In *Conference on Neural Information Processing Systems (NIPS) – Keynote*, Long Beach, California.
- Caroline Criado-Perez. 2019. *Invisible Women: Exposing Data Bias in a World Designed for Men*. Penguin Random House, London, UK.
- Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019a. MuST-C: a Multilingual Speech Translation Corpus. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, page 2012–2017, Minneapolis, Minnesota, June.
- Mattia A. Di Gangi, Matteo Negri, Roldano Cattoni, Dessi Roberto, and Marco Turchi. 2019b. Enhancing transformer for end-to-end speech-to-text translation. In *Machine Translation Summit XVII*, pages 21–31, Dublin, Ireland, August. European Association for Machine Translation.
- Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2019c. One-To-Many Multilingual End-to-end Speech Translation. In *Proceedings of the 2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 585–592, Sentosa, Singapore, December.
- Mostafa Elaraby, Ahmed Y. Tawfik, Mahmoud Khaled, Hany Hassan, and Aly Osama. 2018. Gender-Aware Spoken Language Translation Applied to English-Arabic. In *Proceedings of the 2nd International Conference on Natural Language and Speech Processing (ICNLSP)*, pages 1–6, Algiers, Algeria.

- 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, August. Association for Computational Linguistics.
- Marco Gaido, Mattia A. Di Gangi, Matteo Negri, and Marco Turchi. 2020. End-to-End Speech-Translation with Knowledge Distillation: FBK@IWSLT2020. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 80–88, Online, July. Association for Computational Linguistics.
- Mahault Garnerin, Solange Rossato, and Laurent Besacier. 2020. Gender Representation in Open Source Speech Resources. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 6599–6605, Marseille, France, May. European Language Resources Association.
- GLAAD. 2007. Media Reference Guide - Transgender.
- Marlis Hellinger and Hadumond Bußman. 2001. *Gender across Languages*. John Benjamins Publishing, Amsterdam, The Netherlands.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. 2018. TED-LIUM 3: Twice as Much Data and Corpus Repartition for Experiments on Speaker Adaptation. In *Proceedings of the Speech and Computer - 20th International Conference (SPECOM)*, pages 198–208, Leipzig, Germany, September. Springer International Publishing.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the Knowledge in a Neural Network. In *Proceedings of NIPS Deep Learning and Representation Learning Workshop*, Montréal, Canada, December.
- Charles F. Hockett. 1958. *A Course in Modern Linguistics*. Macmillan, New York, NY, US.
- Dirk Hovy and Anders Søgaard. 2015. Tagging Performance Correlates with Author Age. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 483–488, Beijing, China, July. 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, August. Association for Computational Linguistics.
- Hirofumi Inaguma, Kevin Duh, Tatsuya Kawahara, and Shinji Watanabe. 2019. Multilingual End-to-End Speech Translation. In *Proceedings of the 2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 570–577, December.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Giménez. Adrià, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. Europarl-ST: A Multilingual Corpus For Speech Translation Of Parliamentary Debates. In *Proceedings of 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233, Barcelona, Spain, May.
- Ye Jia, Melvin Johnson, Wolfgang Macherey, Ron J. Weiss, Yuan Cao, Chung-Cheng Chiu, Naveen Ari, Stella Laurenzo, and Yonghui Wu. 2019. Leveraging Weakly Supervised Data to Improve End-to-End Speech-to-Text Translation. In *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7180–7184, Brighton, UK, May.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. Google’s Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation. *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Brian Larson. 2017. Gender as a variable in Natural-PLanguage Processing: Ethical considerations. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 1–11, Valencia, Spain, April. Association for Computational Linguistics.
- Enora Lessinger. 2020. The Challenges of Translating Gender in UN texts. In Luise von Flotow and Hala Kamal, editors, *The Routledge Handbook of Translation, Feminism and Gender*. Routledge, New York, NY, USA.
- Roger J. R. Levesque, 2011. *Sex Roles and Gender Roles*, pages 2622–2623. Springer New York, New York, NY.
- Sarah Ita Levitan, Taniya Mishra, and Srinivas Bangalore. 2016. Automatic identification of gender from speech. In *In Proceedings of Speech Prosody 2016*, pages 84–88, Boston, Massachusetts, May-June.

- Yuchen Liu, Hao Xiong, Jiajun Zhang, Zhongjun He, Hua Wu, Haifeng Wang, and Chengqing Zong. 2019. End-to-End Speech Translation with Knowledge Distillation. In *Proceedings of Interspeech 2019*, pages 1128–1132, Graz, Austria, September.
- Michael McCloskey and Neil J. Cohen. 1989. Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem. *The Psychology of Learning and Motivation*, 24:104–169.
- Amit Moryossef, Roei Aharoni, and Yoav Goldberg. 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, Florence, Italy, August. Association for Computational Linguistics.
- Thai-Son Nguyen, Sebastian Stueker, Jan Niehues, and Alex Waibel. 2020. Improving Sequence-to-sequence Speech Recognition Training with On-the-fly Data Augmentation. In *Proceedings of the 2020 International Conference on Acoustics, Speech, and Signal Processing – IEEE-ICASSP-2020*, Barcelona, Spain, May.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an ASR Corpus Based on Public Domain Audio Books. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210, Brisbane, Australia, April.
- Daniel S. Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D. Cubuk, and Quoc V. Le. 2019. SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition. In *Proceedings of Interspeech 2019*, pages 2613–2617, Graz, Austria, September.
- Tomasz Potapczyk and Pawel Przybycz. 2020. SRPOL’s System for the IWSLT 2020 End-to-End Speech Translation Task. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pages 89–94, Online, July. Association for Computational Linguistics.
- Christina Richards, Walter P. Bouman, Leighton Seal, Meg John Barker, Timo O. Nieder, and Guy T’Sjoen. 2016. Non-binary or Genderqueer Genders. *International Review of Psychiatry*, 28(1):95–102.
- Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. 2018. How2: A Large-scale Dataset For Multimodal Language Understanding. In *Proceedings of Visually Grounded Interaction and Language (ViGIL)*, Montréal, Canada, December. Neural Information Processing Society (NeurIPS).
- 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, July. Association for Computational Linguistics.
- Kristen Schilt and Lauren Westbrook. 2009. ”Gender Normals,” Transgender People, and the Social Maintenance of Heterosexuality. *Gender & Society*, 23(4). 440-464.
- Deven S. Shah, Hansen A. Schwartz, and Dirk Hovy. 2020. Predictive Biases in Natural Language Processing Models: A Conceptual Framework and Overview. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5248–5264, Online, July. 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, July. Association for Computational Linguistics.
- Frederick W. M. Stentiford and Martin G. Steer. 1988. Machine Translation of Speech. *British Telecom Technology Journal*, 6(2):116–122.
- Susan Stryker. 2008. Transgender history, homonormativity, and disciplinarity. *Radical History Review*, 2008(100):145–157.
- Tony Sun, Andrew Gaut, Shirlyn Tang, Yuxin Huang, Mai ElSherief, Jieyu Zhao, Diba Mirza, Elizabeth Belding, Kai-Wei Chang, and William Yang Wang. 2019. Mitigating Gender Bias in Natural Language Processing: Literature Review. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1630–1640, Florence, Italy, July. Association for Computational Linguistics.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the Inception Architecture for Computer Vision. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, Las Vegas, Nevada, United States, June.
- Rachael Tatman. 2017. Gender and dialect bias in YouTube’s automatic captions. In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 53–59, Valencia, Spain, April. Association for Computational Linguistics.

- Jörg Tiedemann. 2016. Opus – parallel corpora for everyone. *Baltic Journal of Modern Computing*, page 384. Special Issue: Proceedings of the 19th Annual Conference of the European Association of Machine Translation (EAMT).
- Eva Vanmassenhove, Christian Hardmeier, and Andy Way. 2018. Getting Gender Right in Neural Machine Translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3003–3008, Brussels, Belgium, October–November. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Proceedings of NIPS 2017*, pages 5998–6008, Long Beach, California, December. NIPS.
- Alex Waibel, Ajay N. Jain, Arthur E. McNair, Hiroaki Saito, Alexander G. Hauptmann, and Joe Tebelskis. 1991. Janus: A speech-to-speech translation system using connectionist and symbolic processing strategies. In *Proceedings of the International Conference on Acoustics, Speech and Signal Processing, ICASSP 1991*, pages 793–796, Toronto, Canada, May 14–17.
- Ron J. Weiss, Jan Chorowski, Navdeep Jaitly, Yonghui Wu, and Zhifeng Chen. 2017. Sequence-to-Sequence Models Can Directly Translate Foreign Speech. In *Proceedings of Interspeech 2017*, pages 2625–2629, Stockholm, Sweden, August.
- Lal Zimman. 2020. Transgender language, transgender moment: Toward a trans linguistics. In Kira Hall and Rusty Barrett, editors, *The Oxford Handbook of Language and Sexuality*.