

Gender Disambiguation in Machine Translation: Diagnostic Evaluation in Decoder-Only Architectures

Chiara Manna Hosein Mohebbi Afra Alishahi
Frédéric Blain Eva Vanmassenhove

Research Center for Cognitive Science and Artificial Intelligence
Tilburg School of Humanities and Digital Sciences
Tilburg University, Netherlands

{c.manna, h.mohebbi, a.alishahi, f.l.g.blain, e.o.j.vanmassenhove}@tilburguniversity.edu

Abstract

While Large Language Models achieve state-of-the-art results across a wide range of NLP tasks, they remain prone to systematic biases. Among these, gender bias is particularly salient in MT, due to systematic differences across languages in whether and how gender is marked. As a result, translation often requires disambiguating implicit source signals into explicit gender-marked forms. In this context, standard benchmarks may capture broad disparities but fail to reflect the full complexity of gender bias in modern MT. In this paper, we extend recent frameworks on bias evaluation by: (i) introducing a novel measure coined "Prior Bias", capturing a model's default gender assumptions, and (ii) applying the framework to decoder-only MT models. Our results show that, despite their scale and state-of-the-art status, decoder-only models do not generally outperform encoder-decoder architectures on gender-specific metrics; however, post-training (e.g., instruction tuning) not only improves contextual awareness but also reduces the masculine Prior Bias. The code and supplementary data used in this work are available at github.com/chiamanna/gender-bias-diagnostics-mt.

Keywords: machine translation, gender bias, large language models, evaluation framework, diagnostic evaluation, prior bias, attention, instruction tuning

1. Introduction

Recent years have seen remarkable advancements in the field of Machine Translation (MT), marked by the shift from statistical to Neural MT, and the introduction of Transformers (Vaswani et al., 2017), which have since become the backbone of most NLP systems. While early implementations relied on an encoder-decoder design (Tiedemann and Thottingal, 2020; Liu et al., 2020; Costa-jussà et al., 2022; Raffel et al., 2020), current state-of-the-art systems are predominantly decoder-only Large Language Models (LLMs) (OpenAI, 2023; Touvron et al., 2023; Gemma Team et al., 2024; Jiang et al., 2024), pretrained on massive text corpora and often fine-tuned for specific applications (Bommasani et al., 2021). Their success largely stems from the ability to build contextualized representations and capture long-range dependencies (Vaswani et al., 2017; Tay et al., 2020), which has led to substantial improvements in translation quality and established LLMs as the new state-of-the-art in MT (Kocmi et al., 2023, 2024; Deutsch et al., 2025). At the same time, awareness of their limitations has grown in parallel (Kirk et al., 2021; Bender et al., 2021; Navigli et al., 2023; Gallegos et al., 2024).

INPUT
The cook prepared a soup for the housekeeper because he helped clean the room.
OUTPUT
Il cuoco ha preparato una zuppa per la governante perché ha aiutato a pulire la stanza.
INPUT
The farmer offered apples to the housekeeper because she had too many of them.
OUTPUT
Il contadino offrì mele alla governante perché ne aveva troppe.

Figure 1: EN→IT WinoMT examples, where translations¹ default to *la governante* (fem.) and *il contadino* (masc.), regardless of the gender cue (*he/she*).

Their growing scale and complexity make it increasingly difficult to trace back decisions or explain systematic errors. Moreover, because they are pretrained on largely uncurated corpora, with little transparency about data composition, it is nearly impossible to ensure high-quality training data. This

¹Generated by ChatGPT on September 29th, 2025.

lack of control raises concerns about the biases that such models inherit, propagate, and potentially amplify, as they become increasingly integrated into one’s everyday life (Bansal, 2022).

Despite extensive research efforts on fairness and bias in NLP (Sun et al., 2019; Costa-jussà, 2019; Blodgett et al., 2020; Stanczak and Augenstein, 2021), one form of bias that remains particularly salient in MT is gender bias, due to systematic differences across languages in whether and how gender is marked (Ackerman, 2019; Cao and Daumé III, 2020): while English primarily marks gender through pronouns, morphologically richer languages such as Italian or German require gender agreement across multiple parts of speech (Stahlberg et al., 2007). This often forces implicit source information to become explicit in the target (Vanmassenhove et al., 2018; Vanmassenhove, 2024). In principle, contextual cues such as gendered pronouns could guide models’ decisions, yet in practice, these are often disregarded in favor of stereotypical associations. Figure 1 illustrates this through two examples from the WinoMT dataset (Stanovsky et al., 2019), where the corresponding translations overlook the contextual cues available: *housekeeper* is realized in the feminine form despite the masculine pronoun *he*, and *farmer* in a masculine one despite the feminine pronoun *she*. Such errors not only reproduce gender inequalities but also risk reinforcing them on a broader scale (Blodgett et al., 2020).

Standard gender accuracy metrics reveal broad disparities in bias evaluation but offer little insight into whether models actually use available contextual gender cues or simply default to stereotypes. To address this, we previously propose an evaluation framework based on Minimal Pair Accuracy (MPA) (Manna et al., 2025), which measures how consistently models integrate contextual gender cues when translating. Specifically, we consider a translation as “accurate” only if swapping the gender cue (e.g. he/she) yields correct translations in both cases, that is, within a minimal pair.

Building on this idea, we extend the framework along two dimensions. First, we introduce **Prior Bias**, a novel measure to quantify a model’s gender assumption in the absence of explicit cues. To this end, we constructed a Neutral Set, the neutralized extension of WinoMT (Stanovsky et al., 2019) where gendered pronouns are replaced with *they/them* and agreement is adjusted accordingly (e.g. adjectives/verbs). Combined with MPA (Manna et al., 2025), this extension allows us to diagnose whether model errors stem from ignoring contexts, defaulting to a prior bias, or both. Secondly, we further apply the extended framework to **decoder-only models**, evaluating and comparing gender bias in MT across both foundational and

instruction-tuned variants. We find that decoder-only models do not surpass encoder-decoder ones on our gender-specific MT metrics; however, post-training (e.g. instruction tuning) improves contextual cue integration and reduces Prior Bias.

2. Related Work

Over the past decade, research on gender bias in NLP has received increasing attention. In 2014, Schiebinger (2014) exposed the risks of overlooking gender in science and technology. As an early example, they pointed to Machine Translation (MT), noting that automatic translation tools frequently defaulted to masculine pronouns, reflecting their disproportionate frequency in training data. Since then, systematic gender bias has been documented across a wide range of domains, from literature (Hoyle et al., 2019), news (Wevers, 2019), and media (Asr et al., 2021), to word embeddings (Bolukbasi et al., 2016), pretrained language models (Nadeem et al., 2021), and downstream tasks, such as sentiment analysis (Kiritchenko and Mohammad, 2018), coreference resolution (Rudinger et al., 2018; Webster et al., 2018; Zhao et al., 2018), hate speech detection (Park et al., 2018), language generation (Sheng et al., 2020), syntactic analysis (Garimella et al., 2019) and, indeed, MT (Vanmassenhove et al., 2018; Savoldi et al., 2021). The latter remains a particularly challenging setting due to the cross-linguistic nature of the task (Stahlberg et al., 2007; Ackerman, 2019; Cao and Daumé III, 2020). Translating from languages where gender is sparsely marked (e.g., English) into morphologically richer ones (e.g., German, Italian or Hindi) often requires making implicit information explicit in the target (Vanmassenhove et al., 2018; Vanmassenhove, 2024). This added complexity increases the likelihood that models may rely on (stereotypical) statistical associations when resolving gender, as illustrated in Figure 1.

Research on gender bias in MT has primarily focused on analyzing system outputs (Rescigno et al., 2020; Ramesh et al., 2021) and developing bias mitigation strategies, targeting different stages of the translation pipeline. These include rewriting translations into explicitly gendered forms (Vanmassenhove et al., 2018; Moryossef et al., 2019; Habash et al., 2019) or neutral alternatives (Vanmassenhove et al., 2021; Sun et al., 2021), applying embedding-level debiasing methods (Hirasawa and Komachi, 2019; Font and Costa-jussà, 2019), and leveraging counterfactual data augmentation (Zmigrod et al., 2019) or domain adaptation (Saunders and Byrne, 2020). Nonetheless, these approaches have generally proven only partially effective and difficult to scale to modern MT systems (Savoldi et al., 2025). The challenge has

been further amplified by the rise of LLMs, which now underpin the state-of-the-art in MT (Kocmi et al., 2023, 2024; Deutsch et al., 2025), alongside most NLP tasks. Built on the Transformer architecture (Vaswani et al., 2017), these models excel in constructing contextualized token representations by capturing long-range dependencies (Tay et al., 2020), but their opacity raises concerns about inherited and amplified biases (Caliskan et al., 2017; Kirk et al., 2021; Bender et al., 2021; Navigli et al., 2023; Gallegos et al., 2024). Moreover, evidence shows that Transformer-based models for Machine Translation often fail to exploit context effectively. Studies on intra- and inter-sentential context usage (Goindani and Shrivastava, 2021; Voita et al., 2021; Sarti et al., 2024; Mohammed and Niculae, 2024) reveal that models frequently ignore relevant context and/or attend to irrelevant tokens, resulting in systematic errors and biased outputs (Kim et al., 2019; Yin et al., 2021).

Within this landscape, available benchmark datasets, such as the WinoMT (Stanovsky et al., 2019; Bentivogli et al., 2020), have provided controlled settings to expose systematic gaps across systems and language pairs, but are insufficient to capture the full complexity of gender bias in modern, state-of-the-art MT. Standard evaluation frameworks largely operate at the surface level, verifying whether gender markings in the output match the source (Figure 1). However, no insights into the mechanisms underlying model decisions are revealed. Particularly, *we cannot determine whether models rely on contextual cues or default to stereotypes, nor can we capture prior gender assumptions*. These limitations underscore the need for more fine-grained evaluation frameworks that combine evaluative and diagnostic purposes, by quantifying disparities in performance while revealing where errors arise from.

3. Experimental Setup

We investigate whether decoder-only LLMs that underpin the current state-of-the-art in MT integrate contextual gender cues during translation and to what extent their outputs reflect a default (prior) gender preference. To do so, we maintain the two dimensions proposed in the original framework (Manna et al., 2025): an output-level evaluation, which examines model behavior through translation outputs; and an attention-based analysis, which probes internal attention patterns as a proxy for cue integration. This section describes the models, prompting strategies, and data used in our experiments, while the detailed setup for each dimension is presented in its respective section.

3.1. Models

We select three decoder-only models of identical size and architecture (7B parameters, 32 layers, 32 attention heads).

1. **Llama2²** (Touvron et al., 2023): a general-purpose pretrained LLM used here as a baseline to assess gender bias in the absence of MT-specific adaptation.
2. **TowerBase³**: a variant of (1) that underwent **continued pretraining** (Alves et al., 2024) on a mixture of monolingual and parallel corpora, following standard MT training practices.
3. **TowerInstruct⁴**: an extension of (2) refined through **instruction tuning** on translation-specific tasks, achieving GPT-4-level performance on several MT benchmarks (Alves et al., 2024).

We opt for **TowerLLM**, one of the few open-weight decoder-only model suites explicitly tuned for MT, with its instruction-tuned variant (3) consistently outperforming other open alternatives across MT benchmarks and competing with closed-source systems like GPT-4 (OpenAI, 2023), even at the scale of 7B parameters (Alves et al., 2024). Moreover, the TowerLLM model suite enables a direct comparison between a general-purpose LLM, our base model (1), and its translation-tuned variants (2–3) under **controlled post-training conditions**, while retaining a realistic model size for practical MT deployment.

For context, we compare our results to the previously studied encoder-decoder systems, as reported in Manna et al. (2025).

3.2. Prompting

MT-specific prompting strategies vary widely across models and language pairs, and no single template has yet proven consistently optimal (Zhang et al., 2023; Leiter and Eger, 2024). While prior work has predominantly focused on few-shot prompting (Hendy et al., 2023; Vilar et al., 2023), our goal is to isolate model behavior rather than to optimize performance through prompt engineering. We therefore adopt the (zero-shot) prompt format used during TowerInstruct’s instruction tuning (Alves et al., 2024), and provide a plain natural language equivalent for Llama2 and TowerBase, which were not trained with predefined prompt templates. The prompt templates used ensure a consistent and comparable setup across the evaluated models (see Fig. 2).

²huggingface.co/meta-llama/Llama-2-7b-hf

³huggingface.co/Unbabel/TowerBase-7B-v0.1

⁴huggingface.co/Unbabel/TowerInstruct-7B-v0.2

Llama2 / TowerBase

```
Translate the following text from English into Italian.  
English: {source_sentence}  
Italian: {generated_translation}
```

TowerInstruct

```
<|im_start|> user  
Translate the following text from English into Italian.  
English: {source_sentence}  
Italian: <|im_end|>  
<|im_start|> assistant  
{generated_translation}
```

Figure 2: Prompt templates for the EN→IT task. Plain for Llama2 and TowerBase, chat-style for TowerInstruct.

3.3. Data

We rely on the WinoMT challenge set (Stanovsky et al., 2019), which builds on Winogender (Rudinger et al., 2018) and WinoBias (Zhao et al., 2018) as a benchmark for gender bias in MT. Each sentence follows a simple template, containing (i) an ungendered profession noun (the target noun to be translated) and (ii) a gendered pronoun that provides the gold gender cue (see Figure 1). This includes a:

- **Regular set:** 3,888 sentences balanced across masculine and feminine referents.
- **Pro-stereotypical and anti-stereotypical subsets:** 1,584 instances each, pairing professions with gendered pronouns aligning with or contradicting gender-role expectations (e.g., *the housekeeper ... she vs. the housekeeper ... he*).

GENDERED

The **developer** visited the hairdresser because **he** needed to cut **his** hair.

NEUTRALIZED

The **developer** visited the hairdresser because **they** needed to cut **their** hair.

Figure 3: Example illustrating the neutralization process followed to construct the Neutral Set.

Neutral set: Following the approach of Vanmassenhove et al. (2021), we additionally release a neutralized version of WinoMT where gendered pronouns are replaced with the neutral form –

they/them – and any adjective/verb agreement is adjusted accordingly (see Figure 3). This set enables a direct measurement of Prior Bias, *i.e.*, a model’s default gender preference in the absence of explicit gender cues, thereby enriching the original framework.

Following prior work, we focus on English→Italian, a contrastive pair where implicit source gender often must be made explicit in the target via agreement (Stahlberg et al., 2007).

4. Output-Level Analysis

Our output-level analysis aims to answer the diagnostic question: *When contextual gender cues are available, do models rely on them, or do they default to prior gender preferences?* We therefore combine three complementary metrics: (i) Standard (Gender) Accuracy, from the integrated WinoMT pipeline, (ii) Minimal Pair Accuracy (MPA) (Manna et al., 2025), and (iii) **Prior Bias**, our novel measure. Together, these allow us to attribute errors to either neglected cues, default gender assumptions, or both.

4.1. Standard Gender Accuracy

We first rely on the integrated WinoMT evaluation pipeline to measure how often the gender of the translated profession noun matches the gold label, as determined by the source pronoun. Input and target sentences are word-aligned using `fast_align`⁵, and the gender of the aligned target term is extracted using `spacy`’s⁶ morphological analyzer. When no gender is detected, the sentence is assigned an UNKNOWN label. Following the original pipeline, such cases are treated as errors, *i.e.*, inaccurate gender realizations, and therefore penalize a model’s accuracy score.

Table 1 reports standard (gender) accuracy metrics – covering overall performance, as well as separate results for masculine and feminine referents – across the regular, pro-, and anti-stereotypical subsets, for both our decoder-only models and the encoder-decoder systems previously studied in Manna et al. (2025).

Two general tendencies hold across architectures: (i) models perform consistently better on masculine referents, and (ii) accuracy is higher in stereotypical contexts where the assigned gender aligns with societal expectations. However, what stands out is that decoder-only models do not generally outperform encoder-decoder systems on gender-related accuracy metrics, regardless of the larger size and architecture.

⁵github.com/clab/fast_align

⁶spacy.io

Moreover, no improvement is observed between Llama2 and TowerBase, despite the latter having been reported to improve general translation quality (Alves et al., 2024). In fact, TowerBase performs worse than the base model on anti-stereotypical cases. By contrast, TowerInstruct shows a substantial overall gain, which comes with feminine accuracy rising sharply from 23.3% to 52.0%. Yet these improvements are not uniform: while overall and feminine accuracies increase, performance on masculine referents decreases. This trade-off, absent from the previously studied encoder-decoder models (Manna et al., 2025), points to a redistribution rather than a uniform improvement.

Set	Model	Overall	Masc.	Fem.
<i>Decoder-Only Models</i>				
REG	Llama2	48.9%	80.8%	23.3%
	TowerBase	47.4%	78.4%	22.6%
	TowerInstruct	58.6%	72.9%	52.0%
PRO-S	Llama2	64.8%	90.2%	39.5%
	TowerBase	65.2%	88.5%	41.9%
	TowerInstruct	75.9%	86.5%	65.4%
ANTI-S	Llama2	40.1%	70.2%	9.9%
	TowerBase	36.6%	64.7%	8.4%
	TowerInstruct	49.7%	58.4%	40.9%
<i>Encoder-Decoder Baselines (Manna et al. 2025)</i>				
REG	OPUS-MT	42.6%	70.1%	20.6%
	NLLB-200	57.0%	79.6%	41.8%
	mBART	60.9%	83.2%	46.5%
PRO-S	OPUS-MT	55.7%	77.3%	34.1%
	NLLB-200	74.9%	87.4%	62.5%
	mBART	76.6%	92.2%	61.0%
ANTI-S	OPUS-MT	34.2%	59.1%	9.2%
	NLLB-200	47.3%	70.4%	24.2%
	mBART	54.0%	71.9%	35.9%

Table 1: Overall, masculine, and feminine (standard) accuracies on WinoMT for our decoder-only models and previously reported encoder-decoder systems (Manna et al., 2025), on the regular (REG), pro-stereotypical (PRO-S), and anti-stereotypical (ANTI-S) sets.

4.2. Minimal Pair Accuracy (MPA)

Models can output the correct gender form without actually relying on the contextual cue, by defaulting to stereotypical associations. MPA (Manna et al., 2025) addresses this limitation by only "rewarding" context-sensitive predictions rather than simply correct ones. Minimal Pairs (see Fig. 4) consist of two otherwise identical sentences that differ only in the gendered pronouns (*he/she*). A pair is considered accurate, only if the model produces the appropriately gendered translation for the profession *in both* cases, therefore demonstrating it adapted based

on the contextual cue (and thus did not simply rely on a stereotypical default).

PRO-S
The analyst consulted with the librarian because she knows a lot about books.
ANTI-S
The analyst consulted with the librarian because he knows a lot about books.

Figure 4: Minimal pair illustration. The pro-stereotypical case (top) assigns *she* to *librarian*, while the anti-stereotypical case (bottom) assigns *he*. Accurate EN→IT translations must adapt grammatical gender: *la bibliotecaria* (f.) vs. *il bibliotecario* (m.).

Model	MPA	Pro-F	Pro-M
<i>Decoder-Only Models</i>			
Llama2	13.1%	73.4%	26.6%
TowerBase	12.3%	77.4%	22.6%
TowerInstruct	36.7%	50.4%	49.6%
<i>Encoder-Decoder Baselines (Manna et al. 2025)</i>			
OPUS-MT	6.1%	82.3%	17.7%
NLLB-200	30.2%	69.1%	30.9%
mBART	38.5%	61.9%	38.1%

Table 2: Minimal Pair Accuracy (MPA) and its breakdown by gender and stereotype for decoder-only models and encoder-decoder baselines (results reported in Manna et al. (2025)). Pro-F and Pro-M indicate the proportion of accurate minimal pairs where the profession is stereotypically associated with women or men, respectively.

As shown in Table 2, MPA scores are considerably lower than standard Gender Accuracy (Table 1), confirming that models often produce the correct gender form without relying on the contextual cue. Moreover, differences across decoder-only models are primarily driven by the instruction-tuned variant. Llama2 reaches a score of 13.1%, which slightly decreases with TowerBase (12.3%). By contrast, instruction tuning marks a clear shift, with TowerInstruct significantly outperforming the base model. When comparing our results to the previously studied encoder-decoder models (Manna et al., 2025), we observe that these achieve comparable or higher MPA values, indicating that state-of-the-art decoder-only LLMs are not inherently better at context-sensitive gender agreement despite their larger scale and broader pretraining.

We further break down our results by reporting the proportion of accurate minimal pairs in which the profession is stereotypically associated with the

feminine (Pro-F) or masculine gender (Pro-M). Consistent with previous findings (Manna et al., 2025), both Llama2 and TowerBase exhibit a pronounced skew: most accurate pairs involve stereotypically feminine professions, whereas stereotypically masculine ones rarely appear in both gender forms (e.g., *mechanic* is never rendered in its Italian feminine form, *la meccanica*). TowerInstruct, however, diverges from this pattern by achieving an almost perfectly balanced split between Pro-F and Pro-M (50.4% / 49.6%). This pattern supports a shift away from a masculine default and toward more consistent cue integration.

4.3. Prior Bias

To measure default gender preferences directly, we turn to our Neutral Set. We apply the same alignment and morphological analysis pipeline that supports WinoMT (Section 4.1) to extract the realized gender of the target profession.

Since no explicit gender cue is present in this setting, there is no notion of correctness and the observed masculine/feminine distribution reflects the model’s **Prior Bias**. Accordingly, the metric is computed over instances in which a gender realization is detected, thereby excluding UNKNOWN cases.

Standard Accuracy			
Model	Overall	Masculine	Feminine
Llama2	48.9%	80.8%	23.3%
TowerBase	47.4%	78.4%	22.6%
TowerInstruct	58.6%	72.9%	52.0%

Diagnostic Evaluation			
Model	MPA	Prior Bias	
		Masc.	Fem.
Llama2	13.1%	85.3%	14.7%
TowerBase	12.3%	85.6%	14.4%
TowerInstruct	36.7%	75.2%	24.8%

Table 3: Standard Accuracy reports traditional WinoMT metrics (Overall, Masculine, Feminine Accuracy). Diagnostic Evaluation combines (i) Minimal Pair Accuracy (MPA); and (ii) **Prior Bias**, capturing the proportion of masculine vs. feminine gendered forms on the neutralized set, *i.e.*, in the absence of explicit cues.

Table 3 summarizes our results across all evaluated metrics, combining standard WinoMT accuracies with both existing and extended diagnostic measures. While standard Gender Accuracy and MPA already pointed to an overall masculine default, the neutral-set evaluation makes this explicit: both Llama2 and TowerBase overwhelmingly default to the masculine gender (~85%), showing no evidence of a preference shift. By contrast, in

line with its higher MPA and more balanced Pro-F/Pro-M distribution, TowerInstruct reduces this default and generates feminine forms more frequently (~25%). Although a masculine prior remains dominant overall, instruction tuning is here associated with both stronger cue integration and a marked reduction in masculine Prior Bias.

5. Attention-Based Analysis

While output-level metrics reveal whether models correctly resolve gender, they do not show how cues are processed internally. In this section, we examine the models’ internals by analyzing the pattern of self-attention – the core mechanism of context mixing in Transformers (Mohebbi et al., 2023) – to understand how contextual cues are integrated into the representation of the target word during the translation process.

5.1. Setup

For each layer and head, we extract the self-attention weight allocated to the gender cue at the time step in which the target word is generated. Since target words are often split into multiple tokens, we implement a matching strategy to retrieve the full span of subwords for each tokenized target, explicitly including inflected articles (*il, lo, la*, etc.), which are the first locus of gender realization in Italian. This step is language-specific but easily adaptable to other target languages by either modifying or omitting the gendered article component.

Attention scores are then averaged across all subwords, and aggregated over accurate minimal pairs, *i.e.*, cases where the model correctly adapted the target to the pronouns (see Section 4.2). Because these cues are the only explicit gender indicators in the source, restricting the analysis to accurate pairs allows us to identify where – across layers and heads – gender information is encoded in the representation of the target words. To ensure comparability across such aggregated scores, we further restrict the analysis to the minimum number of accurate minimal pairs observed across systems ($n = 195$).

After a preliminary inspection, we observe that attention to masculine pronouns is generally weaker, whereas feminine pronouns trigger stronger attention patterns. This suggests that our decoder-only models, like the encoder-decoder ones analyzed in Manna et al. (2025), must actively integrate contextual information to override a masculine default – aligning with the strong masculine Prior Bias identified in Section 4.3. We therefore focus on **feminine cues in accurate minimal pairs**.

The most pronounced increase in self-attention weights can be observed between Llama2 and TowerBase, whereas the most marked reduction of the inherent masculine Prior Bias only emerges with TowerInstruct (cf. Section 4.3). This suggests that, although continued pretraining on MT-related data enables the base model to better identify and incorporate contextual gender cues into the representation of the target words, these improvements are not directly reflected in the model’s output. In fact, it is only post instruction tuning that the model appears to effectively leverage these learned attention signals to mitigate its masculine Prior Bias. This aligns with Choshen et al. (2019), who observe that post-training practices tend to amplify latent signals already learned by the model, rather than introducing entirely new ones.

5.2.1. Sanity Checks

While attention represents a popular choice to interpret and visualize token interactions across layers and heads in MT research, its reliability as explanatory signal has been widely debated (Jain and Wallace, 2019; Bibal et al., 2022). We therefore complement our analysis with targeted sanity checks.

Prompt Attention Although we keep the translation instruction identical for all models (Section 3.2), the TowerInstruct model includes additional special tokens in its chat template. This raises the question of whether attention to gender cues can be directly compared across models. To check this, we measure the total attention each model assigns to its prompt tokens (see Figure 2), *i.e.*, the sum of the attention weights associated to all prompt tokens when generating the target word. We see a highly consistent value across models ($\sim 0.77-0.78\%$), indicating that our direct comparison is fair as the main context receives equivalent total attention across models.

Model	Target (L15)		Secondary (L15)	
	H10	H25	H10	H25
Llama2	0.11	0.12	0.01	0.02
TowerBase	0.17	0.16	0.02	0.01
TowerInstruct	0.19	0.16	0.04	0.03

Table 4: Attention to the **feminine cue** in Layer 15, Heads 10 and 25 when generating the translation of our target v.s. our secondary (non-coreferent) entity.

Secondary Entity We further test whether models might attend to pronouns generically, rather than to resolve the gender of the coreferent entity. Since each WinoMT sentence includes a secondary (non-coreferent) entity (e.g., *The doctor asked the nurse to help her.*), we annotate their positions and measure the attention the model allocates to the cue

when generating their corresponding translations (see Appendix A for the corresponding heatmaps). We observe that this is consistently low (≤ 0.04), suggesting that models attend to pronouns selectively, only when these are relevant for gender disambiguation. Table 4 reports these values for our prominent heads, Heads 10 and 25 at Layer 15, and compares them with the attention from our target in Figure 5.

6. Conclusion

We extend the original evaluation framework (Manna et al., 2025) by (i) enriching it with an explicit measure of **Prior Bias**, *i.e.*, a model’s default gender preference, and (ii) applying it to **decoder-only models** in both their general-purpose and translation-tuned variants.

We find that decoder-only models, which underpin the current state-of-the-art, do not consistently outperform encoder-decoder systems (Manna et al., 2025). This observation resonates with recent evidence that decoder-only LLMs tend to underperform encoder- or encoder-decoder models on tasks requiring more contextual disambiguation (e.g., word meaning tasks (Qorib et al., 2024)). Nonetheless, similar overall tendencies persist: all models perform better on masculine referents and in stereotypical contexts. Moreover, MPA consistently falls below standard Gender Accuracy, revealing that even when models produce the expected target gender, they do not consistently rely on contextual cues. While all evidence points to a default-to-masculine reasoning process – which aligns with prior work (Jumelet et al., 2019; Van Der Wal et al., 2022) – this tendency is now explicitly captured and quantified through our Prior Bias metric. Together, these results help contextualize systematic gender errors in MT, tracing them to ignored contextual cues and pronounced gender priors.

At the same time, our setup enables a controlled comparison across TowerLLM’s post-training stages. While continued pretraining offers little benefit for gender disambiguation, instruction tuning is associated with stronger cue integration - reflected in higher MPA, improved feminine accuracy, and a marked reduction in masculine Prior Bias. Ultimately, our framework provides a diagnostic tool to better understand how gender is encoded and resolved in translation models, and to inform interventions that can improve model behavior. Where traditional gender bias mitigation approaches remain difficult to scale (Savoldi et al., 2025), instruction tuning offers a promising direction towards more contextually aware and gender-fair translation behavior, especially considering the persistent challenge of ensuring unbiased, high-quality training data in large-scale MT and LLMs.

7. Limitations

While our work strengthens the WinoMT framework (Stanovsky et al., 2019), it also inherits some of its limitations. First, the analysis remains constrained to a binary gender distinction, thereby excluding non-binary and gender-neutral alternatives, which are of social and linguistic relevance but remain underrepresented in most current NLP resources. Developing datasets and extending evaluation frameworks beyond binary gender thus remains an important direction for future work.

Second, the evaluation pipeline relies on two automatic processing steps—`fast_align` and `spaCy`—which inevitably introduce some degree of noise into the analysis. In addition to erroneous gender assignments, there are instances in which no gender can be (automatically) detected, resulting in an UNKNOWN label. These may arise from incorrect translations (e.g., incoherent output or omission of the target entity), alignment errors, or inherent limitations of the morphological analyzer. Nonetheless, such cases are rarely explicitly reported, as the pipeline is often applied without further inspection.

Model	Unknown (%)
Llama2	9.4%
TowerBase	10.3%
TowerInstruct	12.8%

Table 5: Percentage of UNKNOWN gender assignments for each model on the regular WinoMT set, *i.e.*, instances where the automatic pipeline failed to detect a gender.

To quantify this phenomenon, we report the proportion of UNKNOWN cases in our analysis (Table 5). Although these are broadly comparable across models and do not alter our qualitative conclusions, even under the conservative assumption that all UNKNOWN cases correspond to correct gender realizations, their frequency deserves attention and may reflect both model-specific behavior and language-specific ambiguity. While this is beyond the scope of this work, a more fine-grained analysis to disentangle genuine gender errors from pipeline-specific or cross-linguistic artifacts is left for future investigation.

Beyond dataset- and evaluation-related considerations, we acknowledge that the attention-based analysis provides correlational rather than causal evidence. Complementing it with causal inter-pretability techniques would allow for a more direct assessment of the role of gender cues in translation decisions. Finally, while instruction tuning emerges as a promising avenue for gender bias mitigation in MT, it remains unclear which specific aspects of the process drive the observed improvements. Further investigation is needed to understand how

post-training practices can contribute to the development of more transparent bias mitigation strategies.

8. Bias Statement

We understand gender bias in MT as a model’s systematic tendency to rely on learned statistical regularities rather than consistently relying on contextual information for gender disambiguation, as previously defined (Manna et al., 2025). Our analysis is restricted to non-ambiguous contexts, in which gender is explicitly signaled in the source – typically through gendered pronouns referring to human entities – representing one specific subtype of gender bias. In these cases, translations often appear to be driven by prior gender associations rather than contextual evidence. This can appear either as stereotypical bias, where models tend to only produce feminine translations when the occupation is socially associated with women (e.g., *housekeeper* becomes *la governante*), or as a masculine default, where models revert to masculine forms regardless of the cue, reinforcing a male-as-norm pattern (Danesi, 2014). Such behaviors can lead to both representational harm, by perpetuating traditional gender roles, and allocational harm, by systematically underrepresenting women in male-dominated professions (Blodgett et al., 2020).

Because the WinoMT dataset (Stanovsky et al., 2019) encodes gender as a binary variable, and its underlying linguistic tools follow the same assumption, our analysis is restricted to masculine and feminine gender realizations. We acknowledge that this is a major limitation, as gender exists along a spectrum rather than a binary divide. However, there is still no broadly adopted methodology to evaluate non-binary gender bias in MT. Developing evaluation frameworks that better reflect gender diversity and inclusivity remains an important direction for future research.

9. Bibliographical References

- Lauren Ackerman. 2019. [Syntactic and cognitive issues in investigating gendered coreference](#). *Glossa: a journal of general linguistics*, 4(1).
- Duarte M. Alves, José Pombal, Nuno M. Guerreiro, Pedro H. Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G.C. de Souza, and André F.T. Martins. 2024. [Tower: An open multilingual large language model for translation-related tasks](#). *arXiv preprint arXiv:2402.17733*.

- Fatemeh Torabi Asr, Mohammad Mazraeh, Alexandre Lopes, Vagrant Gautam, Junette Gonzales, Prashanth Rao, and Maite Taboada. 2021. [The gender gap tracker: Using natural language processing to measure gender bias in media](#). *PLOS ONE*, 16(1):e0245533.
- Rajas Bansal. 2022. A survey on bias and fairness in natural language processing. *arXiv preprint arXiv:2204.09591*.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, pages 610–623.
- 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. Association for Computational Linguistics.
- Adrien Bibal, Rémi Cardon, David Alfter, Rodrigo Wilkens, Xiaou Wang, Thomas François, and Patrick Watrin. 2022. [Is attention explanation? an introduction to the debate](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3889–3900, Dublin, Ireland. Association for Computational Linguistics.
- Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. [Language \(technology\) is power: A critical survey of “bias” in NLP](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5454–5476, Online. Association for Computational Linguistics.
- Tolga Bolukbasi, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam T. Kalai. 2016. [Man is to computer programmer as woman is to homemaker? debiasing word embeddings](#). In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avnika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Nieves, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*. Version 3, revised 12 July 2022.
- Aylin Caliskan, Joanna J. Bryson, and Arvind Narayanan. 2017. [Semantics derived automatically from language corpora contain human-like biases](#). *Science*, 356(6334):183–186.
- Yang Trista 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. Association for Computational Linguistics.
- Leshem Choshen, Lior Fox, Zohar Aizenbud, and Omri Abend. 2019. [On the weaknesses of reinforcement learning for neural machine translation](#). *ArXiv*, abs/1907.01752.
- Marta R Costa-jussà. 2019. An analysis of Gender Bias studies in Natural Language Processing. *Nature Machine Intelligence*, 1(11):495–496.
- Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barrault, Gabriel Mejia Gonzalez, Prangthip

- Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#).
- Marcel Danesi. 2014. *Dictionary of Media and Communications*. Routledge.
- Daniel Deutsch, Eleftheria Briakou, Isaac Caswell, Mara Finkelstein, Rebecca Galor, Juraj Juraska, Geza Kovacs, Alison Lui, Ricardo Rei, Jason Riesa, Shruti Rijhwani, Parker Riley, Elizabeth Salesky, Firas Trabelsi, Stephanie Winkler, Biao Zhang, and Markus Freitag. 2025. [Wmt24++: Expanding the language coverage of wmt24 to 55 languages & dialects](#). *arXiv preprint arXiv:2502.12404*.
- 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.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*, 50(3):1097–1179.
- Aparna Garimella, Carmen Banea, Dirk Hovy, and Rada Mihalcea. 2019. [Women’s syntactic resilience and men’s grammatical luck: Gender-bias in part-of-speech tagging and dependency parsing](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3493–3498, Florence, Italy. Association for Computational Linguistics.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Amélie Héliou, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikuła, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L. Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimenko, Tom Hennigan, Vlad Feinberg, Wojciech Stokowiec, Yu-hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeff Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, and Douglas Eck. 2024. [Gemma: Open models based on gemini research and technology](#). *arXiv preprint arXiv:2403.08295*.
- Akshay Goindani and Manish Shrivastava. 2021. [A dynamic head importance computation mechanism for neural machine translation](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 454–462, Held Online. IN-COMA Ltd.
- Nizar Habash, Houda Bouamor, and Christine Chung. 2019. Automatic gender identification and reinflection in arabic. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 155–165.
- Mohamed Abdelrehim Hendy, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. How good are gpt models at machine translation? a comprehensive evaluation. *arXiv preprint arXiv:2302.09210*.
- Tosho Hirasawa and Mamoru Komachi. 2019. Debiasing word embeddings improves multimodal machine translation. In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 32–42.
- Alexander Miserlis Hoyle, Lawrence Wolf-Sonkin, Hanna Wallach, Isabelle Augenstein, and Ryan Cotterell. 2019. [Unsupervised discovery of gendered language through latent-variable modeling](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1706–1716, Florence, Italy. Association for Computational Linguistics.

- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). 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 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th ophile Gerv t, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2024. [Mixtral of experts](#). *arXiv preprint arXiv:2401.04088*.
- Jaap Jumelet, Willem Zuidema, and Dieuwke Hupkes. 2019. [Analysing neural language models: Contextual decomposition reveals default reasoning in number and gender assignment](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 1–11, Hong Kong, China. Association for Computational Linguistics.
- Yunsu Kim, Duc Thanh Tran, and Hermann Ney. 2019. [When and why is document-level context useful in neural machine translation?](#) In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pages 24–34, Hong Kong, China. Association for Computational Linguistics.
- Svetlana Kiritchenko and Saif Mohammad. 2018. [Examining gender and race bias in two hundred sentiment analysis systems](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 43–53, New Orleans, Louisiana. Association for Computational Linguistics.
- Hannah Rose Kirk, Yennie Jun, Filippo Volpin, Haider Iqbal, Elias Benussi, Frederic Dreyer, Aleksandar Shtedritski, and Yuki Asano. 2021. [Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models](#). In *Advances in Neural Information Processing Systems 34 (NeurIPS 2021)*.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondr ej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Benjamin Marie, Christof Monz, Kenton Murray, Masaaki Nagata, Martin Popel, Maja Popovi c, Mariya Shmatova, Steinh or Steingr msson, and Vil m Zouhar. 2024. [Findings of the WMT24 general machine translation shared task: The LLM era is here but MT is not solved yet](#). In *Proceedings of the Ninth Conference on Machine Translation*, pages 1–46, Miami, Florida, USA. Association for Computational Linguistics.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondr ej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Philipp Koehn, Benjamin Marie, Christof Monz, Makoto Morishita, Kenton Murray, Masaaki Nagata, Toshiaki Nakazawa, Martin Popel, Maja Popovi c, Mariya Shmatova, and Jun Suzuki. 2023. [Findings of the 2023 conference on machine translation \(WMT23\): LLMs are here but not quite there yet](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1–42, Singapore. Association for Computational Linguistics.
- Christoph Leiter and Steffen Eger. 2024. [Prexme! large scale prompt exploration of open source llms for machine translation and summarization evaluation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 11481–11506.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#). *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Chiara Manna, Afra Alishahi, Fr d eric Blain, and Eva Vanmassenhove. 2025. [Are we paying attention to her? investigating gender disambiguation and attention in machine translation](#). In *Proceedings of the 3rd Workshop on Gender-Inclusive Translation Technologies (GIT 2025)*, pages 1–16, Geneva, Switzerland. European Association for Machine Translation.
- Wafaa Mohammed and Vlad Niculae. 2024. [On measuring context utilization in document-level MT systems](#). In *Findings of the Association for Computational Linguistics: EACL 2024*, pages 1633–1643, St. Julian’s, Malta. Association for Computational Linguistics.
- Hosein Mohebbi, Willem Zuidema, Grzegorz Chrupa a, and Afra Alishahi. 2023. [Quantifying context mixing in transformers](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages

- 3378–3400, Dubrovnik, Croatia. Association for Computational Linguistics.
- 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.
- Moin Nadeem, Anna Bethke, and Siva Reddy. 2021. [StereoSet: Measuring stereotypical bias in pre-trained language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5356–5371, Online. Association for Computational Linguistics.
- Roberto Navigli, Simone Conia, and Björn Ross. 2023. [Biases in large language models: Origins, inventory and discussion](#). *Journal of Data and Information Quality*, 15(2):1–21.
- OpenAI. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Ji Ho Park, Jamin Shin, and Pascale Fung. 2018. [Reducing gender bias in abusive language detection](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2799–2804, Brussels, Belgium. Association for Computational Linguistics.
- Muhammad Reza Qorib, Geonsik Moon, and Hwee Tou Ng. 2024. Are decoder-only language models better than encoder-only language models in understanding word meaning? In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 16339–16347.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Krithika Ramesh, Gauri Gupta, and Sanjay Singh. 2021. Evaluating gender bias in hindi-english machine translation. In *Proceedings of the 3rd Workshop on Gender Bias in Natural Language Processing*, pages 16–23.
- Argentina Anna Rescigno, Eva Vanmassenhove, Johanna Monti, and Andy Way. 2020. A case study of natural gender phenomena in translation a comparison of google translate, bing microsoft translator and deepl for english to italian, french and spanish. *Computational Linguistics CLIC-it 2020*, page 359.
- Rachel Rudinger, Jason Naradowsky, Brian Leonard, and Benjamin Van Durme. 2018. [Gender bias in coreference resolution](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 8–14, New Orleans, Louisiana. Association for Computational Linguistics.
- Gabriele Sarti, Grzegorz Chrupała, Malvina Nissim, and Arianna Bisazza. 2024. [Quantifying the plausibility of context reliance in neural machine translation](#). In *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*.
- 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.
- Beatrice Savoldi, Jasmijn Bastings, Luisa Bentivogli, and Eva Vanmassenhove. 2025. A decade of gender bias in machine translation. *Patterns*, 6(6).
- Beatrice Savoldi, Marco Gaido, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. Gender bias in machine translation. *Transactions of the Association for Computational Linguistics*, 9:845–874.
- Londa Schiebinger. 2014. [Scientific research must take gender into account](#). *Nature*, 507(7490):9–9.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2020. [Towards Controllable Biases in Language Generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3239–3254, Online. Association for Computational Linguistics.
- Dagmar Stahlberg, F. Braun, L. Irmen, and Sabine Sczesny. 2007. Representation of the sexes in language. *Social Communication*, pages 163–187.
- Karolina Stanczak and Isabelle Augenstein. 2021. A survey on gender bias in natural language processing. *arXiv preprint arXiv:2112.14168*.
- Gabriel Stanovsky, Noah A. Smith, and Luke Zettlemoyer. 2019. [Evaluating gender bias in machine translation](#). *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- 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. Association for Computational Linguistics.
- Tony Sun, Kellie Webster, Apu Shah, William Yang Wang, and Melvin Johnson. 2021. They, them, theirs: Rewriting with gender-neutral english. *arXiv preprint arXiv:2102.06788*.
- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. 2020. [Long range arena: A benchmark for efficient transformers](#). *ArXiv*, abs/2011.04006.
- Jörg Tiedemann and Santhosh Thottingal. 2020. [OPUS-MT – building open translation services for the world](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 479–480, Lisboa, Portugal. European Association for Machine Translation.
- Hug Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric M. Smith, Ranjan Subramanian, Xiaoqing E. Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang K., Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv (Cornell University)*.
- Oskar Van Der Wal, Jaap Jumelet, Katrin Schulz, and Willem Zuidema. 2022. [The birth of bias: A case study on the evolution of gender bias in an English language model](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 75–75. Association for Computational Linguistics.
- Eva Vanmassenhove. 2024. Gender bias in machine translation and the era of large language models. In *Gendered Technology in Translation and Interpreting*, pages 225–252. Routledge.
- Eva Vanmassenhove, Chris Emmerly, and Dimitar Shterionov. 2021. [NeuTral Rewriter: A rule-based and neural approach to automatic rewriting into gender neutral alternatives](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8940–8948, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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. 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 *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George Foster. 2023. [Prompting PaLM for translation: Assessing strategies and performance](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15406–15427, Toronto, Canada. Association for Computational Linguistics.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2021. [Analyzing the source and target contributions to predictions in neural machine translation](#). 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 1126–1140, Online. Association for Computational Linguistics.
- Kellie Webster, Marta Recasens, Vera Axelrod, and Jason Baldridge. 2018. [Mind the GAP: A balanced corpus of gendered ambiguous pronouns](#). *Transactions of the Association for Computational Linguistics*, 6:605–617.
- Melvin Wevers. 2019. [Using word embeddings to examine gender bias in Dutch newspapers, 1950-1990](#). In *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change*, pages 92–97, Florence, Italy. Association for Computational Linguistics.

Kayo Yin, Patrick Fernandes, Danish Pruthi, Aditi Chaudhary, André F. T. Martins, and Graham Neubig. 2021. [Do context-aware translation models pay the right attention?](#) 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 788–801, Online. Association for Computational Linguistics.

Biao Zhang, Barry Haddow, and Alexandra Birch. 2023. [Prompting large language model for machine translation: A case study.](#) In *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*, volume 202, pages 23276–23295.

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

Ran Zmigrod, Sabrina J Mielke, Hanna Wallach, and Ryan Cotterell. 2019. Counterfactual data augmentation for mitigating gender stereotypes in languages with rich morphology. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1651–1661.

