

Filling Gender & Number Gaps in Neural Machine Translation with Black-box Context Injection

Amit Moryossef[†], Roei Aharoni[†], Yoav Goldberg^{†‡}
{first.last}@gmail.com

[†]Bar Ilan University, Ramat Gan, Israel

[‡]Allen Institute for Artificial Intelligence

Abstract

When translating from a language that does not morphologically mark information such as gender and number into a language that does, translation systems must “guess” this missing information, often leading to incorrect translations in the given context. We propose a black-box approach for injecting the missing information to a pre-trained neural machine translation system, allowing to control the morphological variations in the generated translations *without* changing the underlying model or training data. We evaluate our method on an English to Hebrew translation task, and show that it is effective in injecting the gender and number information and that supplying the correct information improves the translation accuracy in up to 2.3 BLEU on a female-speaker test set for a state-of-the-art online black-box system. Finally, we perform a fine-grained syntactic analysis of the generated translations that shows the effectiveness of our method.

1 Introduction

A common way for marking information about gender, number, and case in language is morphology, or the structure of a given word in the language. However, different languages mark such information in different ways – for example, in some languages gender may be marked on the head word of a syntactic dependency relation, while in other languages it is marked on the dependent, on both, or on none of them (Nichols, 1986). This morphological diversity creates a challenge for machine translation, as there are ambiguous cases where more than one correct translation exists for the same source sentence. For example, while the English sentence “I love language” is ambiguous with respect to the gender of the speaker, Hebrew marks verbs

for the gender of their subject and does not allow gender-neutral translation. This allows two possible Hebrew translations – one in a masculine and the other in a feminine form. As a consequence, a sentence-level translator (either human or machine) must commit to the gender of the speaker, adding information that is not present in the source. Without additional context, this choice must be done arbitrarily by relying on language conventions, world knowledge or statistical (stereotypical) knowledge.

Indeed, the English sentence “I work as a doctor” is translated into Hebrew by Google Translate using the masculine verb form *oved*, indicating a male speaker, while “I work as a nurse” is translated with the feminine form *ovedet*, indicating a female speaker (verified on March 2019). While this is still an issue, there have been recent efforts to reduce it for specific language pairs.¹

We present a simple black-box method to influence the interpretation chosen by an NMT system in these ambiguous cases. More concretely, we construct pre-defined textual hints about the gender and number of the speaker and the audience (the interlocutors), which we concatenate to a given input sentence that we would like to translate accordingly. We then show that a black-box NMT system makes the desired morphological decisions according to the given hint, even when no other evidence is available on the source side. While adding those hints results in additional text on the target side, we show that it is simple to remove, leaving only the desired translation.

Our method is appealing as it only requires simple pre-and-post processing of the inputs and outputs, without considering the system internals, or requiring specific annotated data and training procedure as in previous work (Vanmassenhove et al.,

¹blog.google/products/translate/reducing-gender-bias-google-translate/

2018). We show that in spite of its simplicity, it is effective in resolving many of the ambiguities and improves the translation quality in up to 2.3 BLEU when given the correct hints, which may be inferred from text metadata or other sources. Finally, we perform a fine-grained syntactic analysis of the translations generated using our method which shows its effectiveness.

2 Morphological Ambiguity in Translation

Different languages use different morphological features marking different properties on different elements. For example, English marks for number, case, aspect, tense, person, and degree of comparison. However, English does not mark gender on nouns and verbs. Even when a certain property is marked, languages differ in the form and location of the marking (Nichols, 1986). For example, marking can occur on the head of a syntactic dependency construction, on its argument, on both (requiring *agreement*), or on none of them. Translation systems must generate correct target-language morphology as part of the translation process. This requires knowledge of both the source-side and target-side morphology. Current state-of-the-art translation systems do capture many aspects of natural language, including morphology, when a relevant context is available (Dalvi et al., 2017; Bawden et al., 2018), but resort to “guessing” based on the training-data statistics when it is not. Complications arise when different languages convey different kinds of information in their morphological systems. In such cases, a translation system may be required to remove information available in the source sentence, or to add information not available in it, where the latter can be especially tricky.

3 Black-Box Knowledge Injection

Our goal is to supply an NMT system with knowledge regarding the speaker and interlocutor of first-person sentences, in order to produce the desired target-side morphology when the information is not available in the source sentence. The approach we take in the current work is that of *black-box injection*, in which we attempt to inject knowledge to the input in order to influence the output of a trained NMT system, without having access to its internals or its training procedure as proposed by Vanmassenhove et al. (2018).

We are motivated by recent work by Voita et al. (2018) who showed that NMT systems learn to track coreference chains when presented with sufficient discourse context. We conjecture that there are enough sentence-internal pronominal coreference chains appearing in the training data of large-scale NMT systems, such that state-of-the-art NMT systems can and do track sentence-internal coreference. We devise a wrapper method to make use of this coreference tracking ability by introducing artificial antecedents that unambiguously convey the desired gender and number properties of the speaker and audience.

More concretely, a sentence such as “I love you” is ambiguous with respect to the gender of the speaker and the gender and number of the audience. However, sentences such as “I love you, she told him” are unambiguous given the coreference groups {I, she} and {you, him} which determine *I* to be feminine singular and *you* to be masculine singular. We can thus inject the desired information by prefixing a sentence with short generic sentence fragment such as “She told him:” or “She told them that”, relying on the NMT system’s coreference tracking abilities to trigger the correctly marked translation, and then remove the redundant translated prefix from the generated target sentence. We observed that using a parataxis construction (i.e. “she said to him:”) almost exclusively results in target-side parataxis as well (in 99.8% of our examples), making it easy to identify and strip the translated version from the target side. Moreover, because the parataxis construction is grammatically isolated from the rest of the sentence, it can be stripped without requiring additional changes or modification to the rest of the sentence, ensuring grammaticality.

4 Experiments & Results

To demonstrate our method in a black-box setting, we focus our experiments on Google’s machine translation system (GMT), accessed through its Cloud API. To test the method on real-world sentences, we consider a monologue from the stand-up comedy show “Sarah Silverman: A Speck of Dust”. The monologue consists of 1,244 English sentences, all by a female speaker conveyed to a plural, gender-neutral audience. Our parallel corpora consists of the 1,244 English sentences from the transcript, and their corresponding He-

Speaker	Audience	BLEU
Baseline		18.67
He	–	19.2
He	him	19.25
He	her	19.3
He	them	19.5
I	–	19.84
I	them	20.23
She	–	20.8
She	him	20.82
She	her	20.98
She	them	20.97

Table 1: BLEU results on the Silverman dataset

brew translations based on the Hebrew subtitles.² We translate the monologue one sentence at a time through the Google Cloud API. Eyeballing the results suggest that most of the translations use the incorrect, but default, masculine and singular forms for the speaker and the audience, respectively. We expect that by adding the relevant condition of “female speaking to an audience” we will get better translations, affecting both the gender of the speaker and the number of the audience.

To verify this, we experiment with translating the sentences with the following variations: **No Prefix**—The baseline translation as returned by the GMT system. **“He said:”**—Signaling a male speaker. We expect to further skew the system towards masculine forms. **“She said:”**—Signaling a female speaker and unknown audience. As this matches the actual speaker’s gender, we expect an improvement in translation of first-person pronouns and verbs with first-person pronouns as subjects. **“I said to them:”**—Signaling an unknown speaker and plural audience. **“He said to them:”**—Masculine speaker and plural audience. **“She said to them:”**—Female speaker and plural audience—the complete, correct condition. We expect the best translation accuracy on this setup. **“He/she said to him/her”**—Here we set an (incorrect) singular gender-marked audience, to investigate our ability to control the audience morphology.

4.1 Quantitative Results

We compare the different conditions by comparing BLEU (Papineni et al., 2002) with respect

²The data is obtained from www.opensubtitles.org

to the reference Hebrew translations. We use the `multi-bleu.perl` script from the Moses toolkit (Koehn et al., 2007). Table 1 shows BLEU scores for the different prefixes. The numbers match our expectations: Generally, providing an incorrect speaker and/or audience information decreases the BLEU scores, while providing the correct information substantially improves it - we see an increase of up to 2.3 BLEU over the baseline. We note the BLEU score improves in all cases, even when given the wrong gender of either the speaker or the audience. We hypothesise this improvement stems from the addition of the word “said” which hints the model to generate a more “spoken” language which matches the tested scenario. Providing correct information for both speaker and audience usually helps more than providing correct information to either one of them individually. The one outlier is providing “She” for the speaker and “her” for the audience. While this is not the correct scenario, we hypothesise it gives an improvement in BLEU as it further reinforces the female gender in the sentence.

4.2 Qualitative Results

The BLEU score is an indication of how close the automated translation is to the reference translation, but does not tell us what exactly changed concerning the gender and number properties we attempt to control. We perform a finer-grained analysis focusing on the relation between the injected speaker and audience information, and the morphological realizations of the corresponding elements. We parse the translations and the references using a Hebrew dependency parser.³ In addition to the parse structure, the parser also performs morphological analysis and tagging of the individual tokens. We then perform the following analysis.

Speaker’s Gender Effects: We search for first-person singular pronouns with subject case (*ani*, unmarked for gender, corresponding to the English *I*), and consider the gender of its governing verb (or adjectives in copular constructions such as ‘I am nice’). The possible genders are ‘masculine’, ‘feminine’ and ‘both’, where the latter indicates a case where the none-diacriticized written form admits both a masculine and a feminine reading. We expect the gender to match the ones

³<https://www.cs.bgu.ac.il/~yoavg/software/hebparsers/hebdepparser/>




Figure 1: Gender inflection statistics for verbs governed by first-person pronouns.

requested in the prefix.

Interlocutors’ Gender and Number Effects:

We search for second-person pronouns and consider their gender and number. For pronouns in subject position, we also consider the gender and number of their governing verbs (or adjectives in copular constructions). For a singular audience, we expect the gender and number to match the requested ones. For a plural audience, we expect the masculine-plural forms.

Results: Speaker. Figure 1 shows the result for controlling the morphological properties of the speaker (*{he, she, I} said*). It shows the proportion of gender-inflected verbs for the various conditions and the reference. We see that the baseline system severely under-predicts the feminine form of verbs as compared to the reference. The “He said” conditions further decreases the number of feminine verbs, while the “I said” conditions bring it back to the baseline level. Finally, the “She said” prefixes substantially increase the number of feminine-marked verbs, bringing the proportion much closer to that of the reference (though still under-predicting some of the feminine cases).

Results: Audience. The chart in Figure 2 shows the results for controlling the number of the audience (...*to them* vs nothing). It shows the proportion of singular vs. plural second-person pronouns on the various conditions. It shows a similar trend: the baseline system severely under-predicts the plural forms with respect to the reference translation, while adding the “to them” condition brings the proportion much closer to that of the reference.




Figure 2: Number inflection statistics for second-person pronouns.

4.3 Comparison to Vanmassenhove et al. (2018)

Closely related to our work, Vanmassenhove et al. (2018) proposed a method and an English-French test set to evaluate gender-aware translation, based on the Europarl corpus (Koehn, 2005). We evaluate our method (using Google Translate and the given prefixes) on their test set to see whether it is applicable to another language pair and domain. Table 2 shows the results of our approach vs. their published results and the Google Translate baseline. As may be expected, Google Translate outperforms their system as it is trained on a different corpus and may use more complex machine translation models. Using our method improves the BLEU score even further.

	Male	Female
VHW (2018) Baseline	37.58	37.75
VHW (2018) + TAG	38.71	38.97
Google Translate	39.33	39.02
Google Translate + Prefix	39.95	39.95

Table 2: Comparison of our approach (using Google Translate) to Vanmassenhove et al. (2018) on their English-French gender corpus.

4.4 Other Languages

To test our method’s outputs on multiple languages, we run our pre-and post-processing steps with Google Translate using examples we sourced from native speakers of different languages. For every example we have an English sentence and two translations in the corresponding language,

	English Text	Masculine	Feminine
Hebrew	I am nice	ani nehm ^{ad}	ani nehm ^{ada}
Prefix		ani nehmad	ani nehmada
Spanish	I am delighted	Estoy encantado	Estoy encantada
Prefix		Estoy encantado	Estoy encantada
Portuguese	I was called	Eu fui chamado	Eu fui chamada
Prefix		Eu fui chamado	Eu fui chamado
French	I am patient	je suis patient	je suis patiente
Prefix		je suis patient	je suis patiente
Italian	I am beautiful	Sono bello	Sono bella
Prefix		io sono bello	io sono bella
Russian	I wrote a message	Я написал сообщение	Я написала сообщение
Prefix		Я написа л сообщение	Я написа л сообщение
Czech	I gave her the flower	já jsem jí dal květinu	já jsem jí dala květinu
Prefix		Dala jsem jí květinu	Dala jsem jí květinu
Romanian	I am patient	Sunt răbdător	Sunt răbdătoare
Prefix		Sunt răbdător	Sunt răbdătoare
Catalan	I am rich	sóc ric	sóc rica
Prefix		sóc ric	sóc ric
Polish	I am nice	Jestem miły	Jestem miła
Prefix		Jestem miły	Jestem miła

Table 3: Examples of languages where the speaker’s gender changes morphological markings in different languages, and translations using the prefix “He said:” or “She said:” accordingly

one in masculine and one in feminine form. Not all examples are using the same source English sentence as different languages mark different information. Table 3 shows that for these specific examples our method worked on 6/10 of the languages we had examples for, while for 3/10 languages both translations are masculine, and for 1 language both are feminine.

5 Related Work

Rabinovich et al. (2017) showed that given input with author traits like gender, it is possible to retain those traits in Statistical Machine Translation (SMT) models. Grönroos et al. (2017) showed that incorporating morphological analysis in the decoder improves NMT performance for morphologically rich languages. Burlot and Yvon (2017) presented a new protocol for evaluating the morphological competence of MT systems, indicating that current translation systems only manage to capture some morphological phenomena correctly. Regarding the application of constraints in NMT, Sennrich et al. (2016) presented a method for controlling the politeness level in the generated output. Fidler and Goldberg (2017) showed how to guide a neural text generation system towards

style and content parameters like the level of professionalism, subjective/objective, sentiment and others. Tiedemann and Scherrer (2017) showed that incorporating more context when translating subtitles can improve the coherence of the generated translations. Most closely to our work, Vanmassenhove et al. (2018) also addressed the missing gender information by training proprietary models with a gender-indicating-prefix. We differ from this work by treating the problem in a black-box manner, and by addressing additional information like the number of the speaker and the gender and number of the audience.

6 Conclusions

We highlight the problem of translating between languages with different morphological systems, in which the target translation must contain gender and number information that is not available in the source. We propose a method for injecting such information into a pre-trained NMT model in a black-box setting. We demonstrate the effectiveness of this method by showing an improvement of 2.3 BLEU in an English-to-Hebrew translation setting where the speaker and audience gender can be inferred. We also perform a fine-grained syn-

tactic analysis that shows how our method enables to control the morphological realization of first and second-person pronouns, together with verbs and adjectives related to them. In future work we would like to explore automatic generation of the injected context, or the use of cross-sentence context to infer the injected information.

References

- Rachel Bawden, Rico Sennrich, Alexandra Birch, and Barry Haddow. 2018. Evaluating discourse phenomena in neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, pages 1304–1313.
- Franck Burlot and François Yvon. 2017. Evaluating the morphological competence of Machine Translation Systems. In *2nd Conference on Machine Translation (WMT17)*, volume 1, pages 43–55, Copenhagen, Denmark. Association for Computational Linguistics.
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, and Stephan Vogel. 2017. Understanding and improving morphological learning in the neural machine translation decoder. In *IJCNLP*.
- Jessica Fidler and Yoav Goldberg. 2017. Controlling linguistic style aspects in neural language generation. *CoRR*, abs/1707.02633.
- Stig-Arne Grönroos, Sami Virpioja, and Mikko Kurimo. 2017. Extending hybrid word-character neural machine translation with multi-task learning of morphological analysis. In *Proceedings of the Second Conference on Machine Translation*, pages 296–302. Association for Computational Linguistics.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th annual meeting of the ACL on interactive poster and demonstration sessions*, pages 177–180. Association for Computational Linguistics.
- Johanna Nichols. 1986. Head-marking and dependent-marking grammar. *Language*, 62(1):56–119.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.
- Ella Rabinovich, Raj Nath Patel, Shachar Mirkin, Lucia Specia, and Shuly Wintner. 2017. Personalized machine translation: Preserving original author traits. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1074–1084. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 35–40. Association for Computational Linguistics.
- Jörg Tiedemann and Yves Scherrer. 2017. Neural machine translation with extended context. In *Proceedings of the Third Workshop on Discourse in Machine Translation*, pages 82–92. Association for Computational Linguistics.
- Eva Vanmassenhove, Christian Hardmeier, and Andy Way. 2018. Getting gender right in neural mt. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3003–3008, Brussels, Belgium. Association for Computational Linguistics.
- Elena Voita, Pavel Serdyukov, Rico Sennrich, and Ivan Titov. 2018. Context-aware neural machine translation learns anaphora resolution. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1264–1274, Melbourne, Australia. Association for Computational Linguistics.