

Towards Debiasing Translation Artifacts

Koel Dutta Chowdhury¹, Rricha Jalota¹,
Cristina España-Bonet², and Josef van Genabith^{1,2}

¹Saarland University, Saarland Informatics Campus, Germany

²German Research Center for Artificial Intelligence (DFKI)

koel.duttachowdhury@uni-saarland.de

rrja00001@stud.uni-saarland.de

{cristinae, Josef.Van_Genabith}@dfki.de

Abstract

Cross-lingual natural language processing relies on translation, either by humans or machines, at different levels, from translating training data to translating test sets. However, compared to original texts in the same language, translations possess distinct qualities referred to as translationese. Previous research has shown that these translation artifacts influence the performance of a variety of cross-lingual tasks. In this work, we propose a novel approach to reducing translationese by extending an established bias-removal technique. We use the Iterative Null-space Projection (INLP) algorithm, and show by measuring classification accuracy before and after debiasing, that translationese is reduced at both sentence and word level. We evaluate the utility of debiasing translationese on a natural language inference (NLI) task, and show that by reducing this bias, NLI accuracy improves. To the best of our knowledge, this is the first study to debias translationese as represented in latent embedding space.

1 Introduction

“Translationese” refers to features in (professional human or machine) translated text that distinguishes it from non-translated, original text in the same language. Carriers of translationese include lexical and word order choices that are influenced by the source language (Gellerstam, 1986), as well the use of more explicit and standardised constructions (Baker et al., 1993) compared to original text. Translationese has a significant impact in machine translation evaluation. Toral et al. (2018) found that translating source sentences that are already the result of translation are easier to translate than original sentences. Similarly, Edunov et al. (2020) show that back-translation results in large BLEU scores when translationese is on the source side and original text is used as reference. To avoid such artifacts, it is advised to use original source sen-

tences for machine translation evaluation (Zhang and Toral, 2019; Graham et al., 2020). Riley et al. (2020) train sentence-level classifiers to differentiate translationese from original target text, and then use this classifier to tag the training data for an NMT model to produce output that shows fewer translationese effects. Translation-inherited artifacts have been shown to have significant impact on other tasks as well. For example, Singh et al. (2019) show that substituting segments of original training samples by their translations from another language improves performance on natural language inference (NLI) tasks. Clark et al. (2020) introduce a translation-free Question Answering dataset to avoid having inflated gains from translation artifacts in transfer-learning tasks. Artetxe et al. (2020) show that cross-lingual models suffer from induced translation artifacts when evaluated on translated test sets. These examples motivate the need for reducing translation artifacts.

While a number of methods to remove or attenuate human-like biases (e.g., gender, race, etc.) in both static and contextualised word embeddings have recently been proposed (Bolukbasi et al., 2016; Gonen and Goldberg, 2019; Dev and Phillips, 2019; Ravfogel et al., 2020; Liang et al., 2020; Zhou et al., 2021), attenuating and eliminating a more implicit signal like translationese in embeddings has yet to be studied. Translationese signals are complex and multi-faceted and, unlike e.g. gender and profanity, can in general not be captured in terms of simple lists of contrastive word pairs (woman-man, she-he, etc.), but rather manifest as a complex mix of morphological, lexical, syntactic and semantic phenomena. Our study is a first attempt to directly debias translationese encoded as latent representations, based on the recently proposed Iterative Nullspace Projection (INLP) algorithm (Ravfogel et al., 2020).

The main contributions of our work are as follows. (i) We propose to reduce the bias induced

by translation by extending the INLP approach to translationese debiasing at both the word and sentence level, with findings demonstrating that our debiasing approaches can effectively attenuate translationese biases in both static word embedding spaces as well as sentence representations based on contextualised embeddings. (ii) We use the INLP method with a number of neural sentence-level translationese classification architectures (fastText, mBERT, XLM), and propose two alternative methods for detecting explicit translationese bias in word embeddings, and find that after debiasing, the models’ performance on classifying translationese degrades to that of a random classifier. (iii) Finally, by integrating the proposed debiasing method within the NLI task, we show the effect of translation artifacts and their removal on the task.

2 Debiasing Strategies

Much previous research (Bolukbasi et al., 2016; Zhao et al., 2018; Dev and Phillips, 2019; Ravfogel et al., 2020) focuses on eliminating bias in word embeddings. However, existing models use lists of contrastive word pairs (e.g., woman-man, she-he) to detect, capture and mitigate specific biases (e.g., gender, profanity, etc.). While translationese cannot, in general, be captured in simple lists of contrastive word pairs, we do have labeled data at sentence level: translated (translationese) and original sentences. We use this data and the INLP algorithm to directly mitigate translationese at sentence level (Section 3). Additionally, we explore INLP at word level to debias translationese in Section 4. Only a few earlier studies (Dutta Chowdhury et al., 2020, 2021) deal with translationese at the level of word embedding spaces leveraging distances between graph-based representations of original and translationese data. For debiasing word embedding spaces, we adapt an idea from Gonen et al. (2020) to extract lists of pairs of identical words and examine how their use differs in translationese and original data (rather than contrasting word pairs). If a word is used very differently in translated and original data, this is reflected in differences in original and translated word embedding spaces, and is evidence of translationese in the embeddings. Alternatively, we propose a simpler approach that builds on a joint embedding space where words are tagged according to their origin (translationese or original) and without any need for a word list.

The **Iterative Nullspace Projection algorithm**

(Ravfogel et al., 2020) focuses on removing linearly decipherable features from vector representations originally for gender bias mitigation. Given a set of labeled data with data points $X = x_1, \dots, x_n$ and task labels $Y = y_1, \dots, y_n$, we use a standard classification setup with a neural network and a simple classifier τ on top. An encoder h encodes x_k into a representation vector $h(x_k)$ and τ predicts y_k based on $h(x_k)$, i.e., $y_k = \tau(h(x_k))$. Let T be the trait to be mitigated, also known as the *protected attribute*. The goal of the INLP method is to neutralise the ability of the classifier τ to linearly predict T from h . τ is parameterised by a matrix W and trained to predict T from h . Using W , one can collapse the data onto its nullspace $N(W)$ with a projection matrix $P_{N(W)}$. This guarantees $W P_{N(W)} h(T) = 0$, i.e., the information used to classify T is linearly removed. By repeating this process i times until no classifier achieves above-majority accuracy, INLP can neutralise all features that W_i uses for predicting T from \tilde{h} :

$$\tilde{h} := P_{N(W_1)} P_{N(W_2)} \dots P_{N(W_i)} h \quad (1)$$

Details of the implementation are the same as in Ravfogel et al. (2020).¹

3 Translationese in Sentence Embeddings

In our work, we are interested in adapting INLP to study the impact of removing the translationese attributes T from semantic representations, via a binary classification task. Specifically, the binary classifier learns to distinguish between original and translationese sentences. Therefore, in our setup, labels Y correspond to original and translationese and act as protected attributes.

Data. We use the Europarl corpus annotated with translationese information from Amponsah-Kaakyire et al. (2021). We focus on three languages: English (En), German (De) and Spanish (Es). The corpus provides originals in the three languages (L1) and translations into these three languages that come from original texts in the other two (L2). We use the notation L1–L2 to refer to the different sets in Table 1 and Table 2. For example, in En-De, L2 refers to English text translated from German. For each corpus, there is an equal number of translated and original sentences: 42k for

¹Our code is available at <https://github.com/koelddc/Towards-Debiasing-Translation-Artifacts/>

	fastText Text	mBERT CLS	mBERT pool	XLM CLS	de- biased
En-De	0.64	0.73	0.79	0.71	0.50
En-Es	0.71	0.78	0.83	0.77	0.50
De-En	0.68	0.78	0.84	0.77	0.50
De-Es	0.69	0.79	0.86	0.79	0.50
Es-De	0.71	0.77	0.85	0.77	0.50
Es-En	0.72	0.76	0.82	0.77	0.50

Table 1: Sentence Embedding Classification accuracy on original versus translationese using different models. After INLP debiasing, translationese classification reduces to random 50% accuracy in all cases.

De–En, De–Es, En–De, En–Es, Es–De and Es–En. We use 70% of the sentences for training, 15% for development and 15% for testing.

Classifier. We use a logistic classifier on top of sentence embeddings h obtained with 4 models without any additional fine-tuning to the translationese classification task. (i) **fastText** (Joulin et al., 2016): we compute an average of all token vectors in the sentence. (ii) **mBERT_{CLS}** (Devlin et al., 2019): we use the [CLS] token in mBERT as sentence representation. (iii) **mBERT_{pool}** (Devlin et al., 2019): we use mean pooling of mBERT’s contextualised word embeddings. (iv) **XLM_{CLS}** (Conneau et al., 2020): we use the [CLS] token from XLM-RoBERTa.

Results. The first four columns in Table 1 summarise the translationese classification accuracy achieved by the four models. mBERT_{pool} achieves the best performance for all languages, while fastText trails the pack. The final column in Table 1 shows that INLP is close to perfection in removing translationese signals for the linear classifiers, reducing accuracy to a random 50%.

4 Translationese in Word Embeddings

Unlike sentence-level debiasing, word-level debiasing needs a seed translationese direction to obtain a debiased space. This is challenging for translationese as unlike, e.g. gender and profanity, translationese cannot in general be captured in terms of simple contrastive word pairs. In what follows, we introduce two approaches for debiasing translationese at the level of word embeddings.

Stepwise Aligned Space Projection. In order to estimate the seed translationese direction, we derive a list of words (G) used differently in translationese T and original O data using the *usage change* concept from Gonen et al. (2020). The same word used in different data sets (original and

translated) is likely to have different neighbours in the two embedding spaces. We only use words from the intersection of both vocabularies O and T . We compute the score for context change across the embeddings O and T of the two data sets by considering the size of the intersection of two sets where each word in a corpus is represented as its top- k nearest neighbors (NN) in its embedding space:

$$score^k(w) = -|\text{NN}_O^k(w) \cap \text{NN}_T^k(w)| \quad (2)$$

where $\text{NN}_i^k(w)$ is the set of k -NN of word w in embedding space i . The smaller the size of the intersection, the more differently the word is used in the two data sets (and words with the smallest intersection can be seen as indicators of translationese). Given O and T , we collect a ranked list of about 500 words with the smallest intersection as our translationese word list G . G allows us to identify the seed translationese direction for INLP. In our experiments we only consider words attested at least 200 times in the data. Appendix A.3 shows the top 50 elements for all the word lists. In our experiments, we use the translated and original parts of the data described in Section 3 to estimate the word embeddings O and T and use $k=1000$ nearest neighbours in Equation 2. Following Gonen et al. (2020), a large value of k results in large neighbors sets for each word in the two corpora, resulting in a more stable translationese wordlist G .

Next, we create a joint word embedding space \mathcal{J} from the concatenation of the translated and original data, T and O . Since this joint space \mathcal{J} includes both original and translationese signals, we then align the previously unrelated O and T spaces to this embedding space \mathcal{J} , using VecMap (Artetxe et al., 2018), producing aligned spaces \tilde{O} and \tilde{T} , and resulting in an extended single embedding space where T and O are aligned to \mathcal{J} . Next, we compute the translationese direction v of the same word w in the two embeddings spaces, \tilde{O} and \tilde{T} , using,

$$v := \tilde{T}[w] - \tilde{O}[w], \forall w \in G \quad (3)$$

Finally, we compute the similarity of words in \mathcal{J} along the directions v and $-v$, to divide them into two subspaces, translationese and non-translationese, respectively. Using Equation 3, we initialise the INLP algorithm in two ways: (i) **INLP.single**: with a direction vector created from the difference between two aligned spaces for the highest ranked word in G , and (ii) **INLP.avg**: by averaging the differences of all words in G .

	Direct Joint	Stepwise Aligned		de-biased
		INLP.Single	INLP.Avg	
En-De	0.98	0.99	1.00	0.50
En-Es	0.92	1.00	1.00	0.50
De-En	0.96	1.00	0.99	0.50
De-Es	0.91	1.00	1.00	0.50
Es-De	0.93	0.99	0.99	0.50
Es-En	0.95	1.00	1.00	0.50

Table 2: Classification accuracy on original versus translationese with word embeddings using our two approaches, before and after debiasing with INLP.

Direct Joint Space Projection. In this approach, we directly build the embeddings of a specific word w from \mathcal{O} and \mathcal{T} into the same space \mathcal{J}' by annotating w as either w_o and w_t on surface in the \mathcal{O} and \mathcal{T} data. In this way we can easily track and distinguish the two embeddings of the same word w coming from \mathcal{O} and \mathcal{T} data in the same embedding space \mathcal{J}' resulting from the simple concatenation of two datasets. This eliminates the need to compute the translationese direction vector v to group the subspaces and the complexity of maintaining and aligning \mathcal{O} , \mathcal{T} , \mathcal{J} , $\tilde{\mathcal{O}}$ and $\tilde{\mathcal{T}}$ spaces. Figure 1 in Appendix A.2.2 shows the t-SNE projection (Van der Maaten and Hinton, 2008) of the tagged tokens before and after debiasing.

Results. We compare the performance before and after debiasing using our two word-level debiasing methods in Table 2. As expected, debiasing reduces classification accuracy for all language pairs from $\sim 100\%$ to $\sim 50\%$ for both methods.

Translationese Word Lists. The content of our extracted translationese word lists depends on the language (see Appendix A.3). For Es, punctuation is clearly used differently in originals and translations into Es. For De, pronominal adverbs play an important role, especially when translations come from Es. For En, there is no clear trend but, interestingly, only one word in the top-50 list (*indeed*) overlaps with the words with a highest difference in frequency of usage in the original and translationese corpus as analysed for Europarl in Koppel and Ordan (2011). The number of times and the context where a word appears may reflect two different aspects of translationese.

5 Application to NLI

In order to investigate the impact of removing translationese artifacts from the translated data we analyse its impact on the NLI task, where machine translation is used. NLI predicts the relationship

between two sentences, premise and hypothesis, and classifies it into one of the three categories—entailment, contradiction, or neutral. Recently, Artetxe et al. (2020) showed that, in the existing NLI datasets, there exists a significant lexical overlap between the premise and the hypothesis, which is utilised by neural NLI models to make predictions with high accuracies. However, when the premise and hypothesis are paraphrased independently using translation and back-translation, lexical overlap is reduced, negatively impacting the performance of the models.

Below we test whether and if so, to which extent our INLP-based translationese debiasing approach can avoid the observed performance loss in the back-translated NLI task. We generate Back-Translated (BT) NLI dataset by independently translating premise and hypothesis from the Original SNLI data. Then, we train two NLI models, one on Original and one on BT data. To reduce the translationese artifacts resulting from back-translation, we apply the sentence and word embedding debiasing strategies as described in Sections 3 and 4 over the embeddings generated from BT data, and use the resulting debiased-BT embeddings to train a third debiased NLI model in Table 3.

Finally, we consider two scenarios to evaluate translationese debiasing - (i) Symmetric [Sym]: where Original is tested with original test set, and BT and debiased NLI models are tested with BT test data and (ii) Asymmetric [Asym]: where all these models are tested with original test data. We use (i) to see whether debiasing the model trained on translated train-test can bring its performance closer to that of model trained on original train-test data. In (ii), we examine the asymmetry between original test data and BT-NLI training data, and whether our translationese debiasing of BT training data can offset this asymmetry and improve NLI performance. Table 3 shows the classification accuracies with respect to (i) and (ii).

Data. The large-scale SNLI dataset (Stanford Natural Language Inference) (Bowman et al., 2015) contains 570k sentence pairs in the training set manually classified as entailment, contradiction, or neutral. We use a subset of 10% of the training data for our experiments. The development and test data are used as in the original SNLI splits (each containing 10k examples). We generate a back-translated variant of the training and test data using German (De)

Approach	Test Data	SNLI Model		
		Original	Back-translated	debiased
Word-Joint	Sym	67.2±0.1	64.1±0.2	64.7±0.1
	Asym		64.4±0.2	65.1±0.1
Word-Aligned	Sym	67.8±0.1	64.3±0.2	64.8±0.2
	Asym		64.9±0.2	65.3±0.1
Sentence	Sym	65.6±0.0	58.4±0.0	60.1±0.0
	Asym		63.5±0.0	64.9±0.0

Table 3: Test set accuracies (3 runs mean) of our three NLI models in Symmetric and Asymmetric settings.

as the pivot language. For translation, we use the pre-trained models of Facebook-FAIR’s WMT-19 news translation task submission (Ng et al., 2019).

Models. We train three different NLI models, each for original, back-translated and debiased versions of back-translated embeddings. For the word-level setup, we use a single hidden BiLSTM layer followed by a standard feedforward output layer on top of frozen fastText word embeddings for the 3-class NLI classification. The computation of word embeddings and debiasing follows the setup described in Section 4. For sentence-level debiasing, we use the BERT_{pool} method explained in Section 3 with a linear SVC on top to predict the labels.

Results. Table 3 shows results consistent with Artetxe et al. (2020), in that models trained on Original data outperform models trained on BT data in both Sym and Asym scenarios. Table 3 also shows that, after translationese debiasing, classification accuracy on SNLI-debiased improves modestly for all models, with only a minor improvement at word-level, and larger improvement at the sentence level. Overall this may be due to the fact that translationese is a combination of lexical and syntactic phenomena that is better captured at sentence-level. Results in Table 3 suggest that debiasing translation artifacts helps in reducing the asymmetry between translated train and original test set. Therefore, rather than translating the entire test set to match the training set in transfer-learning tasks, debiasing the training set for translation artifacts is a promising direction for future work. Finally, for a complex task such as translationese debiasing, linear intervention alone may not be sufficient. As a result, non-linear guarding approaches need to be investigated further.

6 Conclusion

In this work, we remove translationese artifacts by extending the debiasing INLP approach at both word and sentence level. To the best of our knowl-

edge, this is the first paper that attempts at debiasing sentence and word embeddings for translationese. We introduce two techniques for debiasing translationese at the word level: one (Stepwise Aligned Subspaces) is akin to the subspace construction approach of gender-debiasing proposed by Bolukbasi et al. (2016) and Ravfogel et al. (2020), the second (Direct Joint Subspace) is a simplified approach that operates directly on the joint space without the use of a separate translationese word list and multiple independently computed and subsequently aligned subspaces. Our word-based debiasing study provides a systematic view of translationese biases contained in static embeddings. We also explore translationese debiasing at sentence level embeddings computed from contextualised word embeddings. As expected, the INLP-based linear translationese debiasing results on static word embeddings are as “perfect” as our sentence level results, reducing the performance of a linear translationese classifier on the debiased data to chance, demonstrating that our debiasing strategies effectively attenuate translationese signals in both these spaces.

Further, we evaluate the effects of debiasing translation artifacts on a standard NLI task in two settings. Even though we achieve “perfect” performance for the translationese classification-based debiasing task with INLP, this translates into just modest improvements resulting from INLP-based debiasing translationese in neural machine translation in an NLI task, with slightly better results for sentence than word debiasing. This demonstrates that our debiasing approach is effective in reducing translation artifacts but that there is more to translationese than is visible to a linear classifier.

Finally, we acknowledge that while this study is the first to debias translationese encoded as latent representation in (word and sentence) embedding space, the effect of this on the actual surface form of the generated output is not investigated. We hope to account for this in future work.

7 Acknowledgments

We would like to thank Saadullah Amin for his helpful feedback. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1102 Information Density and Linguistic Encoding.

References

- Kwabena Amponsah-Kaakyire, Daria Pylypenko, Cristina España-Bonet, and Josef van Genabith. 2021. [Do not rely on relay translations: Multilingual parallel direct Europarl](#). In *Proceedings for the First Workshop on Modelling Translation: Translatology in the Digital Age*, pages 1–7, online. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. [A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798, Melbourne, Australia. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2020. [Translation artifacts in cross-lingual transfer learning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7674–7684, Online. Association for Computational Linguistics.
- Mona Baker, Gill Francis, and Elena Tognini-Bonelli. 1993. Corpus linguistics and translation studies: Implications and applications. In *Text and Technology: In Honour of John Sinclair*, page 233–, Netherlands. John Benjamins Publishing Company.
- 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](#). *Advances in Neural Information Processing Systems*, 29:4349–4357.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Sunipa Dev and Jeff Phillips. 2019. [Attenuating bias in word vectors](#). In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 879–887. PMLR.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Koel Dutta Chowdhury, Cristina España-Bonet, and Josef van Genabith. 2020. [Understanding translationese in multi-view embedding spaces](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6056–6062.
- Koel Dutta Chowdhury, Cristina España-Bonet, and Josef van Genabith. 2021. [Tracing source language interference in translation with graph-isomorphism measures](#). In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)*, pages 375–385.
- Sergey Edunov, Myle Ott, Marc’Aurelio Ranzato, and Michael Auli. 2020. [On the evaluation of machine translation systems trained with back-translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2836–2846, Online. Association for Computational Linguistics.
- Martin Gellerstam. 1986. Translationese in Swedish novels translated from English. *Translation studies in Scandinavia*, 1:88–95.
- Hila Gonen and Yoav Goldberg. 2019. [Lipstick on a pig: Debiasing methods cover up systematic gender biases in word embeddings but do not remove them](#). 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 609–614, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hila Gonen, Ganesh Jawahar, Djamé Seddah, and Yoav Goldberg. 2020. [Simple, interpretable and stable method for detecting words with usage change across corpora](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 538–555, Online. Association for Computational Linguistics.
- Yvette Graham, Barry Haddow, and Philipp Koehn. 2020. [Statistical power and translationese in machine translation evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 72–81, Online. Association for Computational Linguistics.

- Armand Joulin, Edouard Grave, Piotr Bojanowski, Matthijs Douze, H erve J egou, and Tomas Mikolov. 2016. Fasttext. zip: Compressing text classification models. *arXiv preprint arXiv:1612.03651*.
- Moshe Koppel and Noam Ordan. 2011. **Translationese and its dialects**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1318–1326, Portland, Oregon, USA. Association for Computational Linguistics.
- Ira Leviant and Roi Reichart. 2015. Separated by an un-common language: Towards judgment language informed vector space modeling. *arXiv preprint arXiv:1508.00106*.
- Sheng Liang, Philipp Dufter, and Hinrich Sch utze. 2020. **Monolingual and multilingual reduction of gender bias in contextualized representations**. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5082–5093, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of machine learning research*, 9(11).
- Nathan Ng, Kyra Yee, Alexei Baevski, Myle Ott, Michael Auli, and Sergey Edunov. 2019. **Facebook FAIR’s WMT19 news translation task submission**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 314–319, Florence, Italy. Association for Computational Linguistics.
- Shauli Ravfogel, Yanai Elazar, Hila Gonen, Michael Twiton, and Yoav Goldberg. 2020. **Null it out: Guarding protected attributes by iterative nullspace projection**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7237–7256, Online. Association for Computational Linguistics.
- Parker Riley, Isaac Caswell, Markus Freitag, and David Grangier. 2020. **Translationese as a language in “multilingual” NMT**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7737–7746, Online. Association for Computational Linguistics.
- Jasdeep Singh, Bryan McCann, Nitish Shirish Keskar, Caiming Xiong, and Richard Socher. 2019. Xlda: Cross-lingual data augmentation for natural language inference and question answering. *arXiv preprint arXiv:1905.11471*.
- Antonio Toral, Sheila Castilho, Ke Hu, and Andy Way. 2018. **Attaining the Unattainable? Reassessing Claims of Human Parity in Neural Machine Translation**. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 113–123, Brussels, Belgium. Association for Computational Linguistics.
- Ivan Vuli c, Simon Baker, Edoardo Maria Ponti, Ulla Petti, Ira Leviant, Kelly Wing, Olga Majewska, Eden Bar, Matt Malone, Thierry Poibeau, Roi Reichart, and Anna Korhonen. 2020. **Multi-SimLex: A large-scale evaluation of multilingual and crosslingual lexical semantic similarity**. *Computational Linguistics*, 46(4):847–897.
- Mike Zhang and Antonio Toral. 2019. **The effect of translationese in machine translation test sets**. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 73–81, Florence, Italy. Association for Computational Linguistics.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2018. **Gender bias in coreference resolution: Evaluation and debiasing methods**. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 15–20, New Orleans, Louisiana. Association for Computational Linguistics.
- Xuhui Zhou, Maarten Sap, Swabha Swayamdipta, Yejin Choi, and Noah Smith. 2021. **Challenges in automated debiasing for toxic language detection**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3143–3155, Online. Association for Computational Linguistics.

A Appendix

A.1 Experimental Setup and Hyperparameters

Each run in Table 1 took around 1.5 hours on a GTX1080Ti GPU. Each classification and debiasing step for SNLI sentence-level experiment in Table 3 took approximately 2 hours on V100-32GB GPU. Other hyperparameter settings are shown in the Table 4.

Model	Hyperparameters
fastText	minCount=5, dim=300
Logistic Regression	warm_start = True, penalty = 'l2', verbose=5, solver="saga", random_state=23, max_iter=7
BiLSTM	hid_dim=300, dropout=0.2, batch_size=32, Adam optimizer with lr=0.0001, epochs=15
SNLI-sentence	nclfs=45, max_iter =1500
SNLI-Aligned (word)	nclfs=34
SNLI-Joint (word)	INLP.single,word=human nclfs=35

Table 4: Hyperparameter settings.

A.2 Debaised Word Representations

A.2.1 Word Analogy Tests

To verify that debiasing does not hurt the quality of the word representations, we estimate the performance of the original and debaised embeddings on the word analogy task using the MultiSIMLEX benchmark (Vulić et al., 2020). As MultiSIMLEX does not cover German, we use German-Simlex from Leviant and Reichart (2015). After debiasing, Spearman’s ρ correlation coefficients show negligible decreases of 0.02 on En-De and En-Es, 0.01 on Es-En and Es-De, 0.3 on De-En and an increase of 0.01 for De-Es.

A.2.2 Visualisation

Figure 1 shows the t-SNE (Van der Maaten and Hinton, 2008) projection of the vectors in Direct Joint Space Projection before and after debiasing with INLP.

De-En	De-Es	En-De	En-Es	Es-De	Es-En
466	510	429	483	547	504

Table 5: Size of translationese word lists created with the usage change algorithm (Gonen et al., 2020).

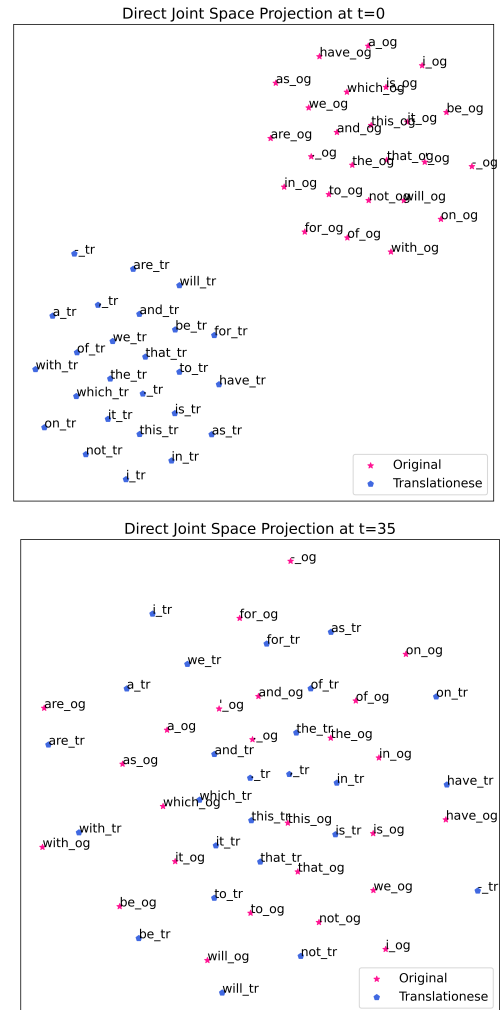


Figure 1: Clustering on Direct Joint Space, before (up-side) and after (down-side) debiasing.

A.3 Translationese Word Lists

The size of the translationese word lists created via the usage change algorithm of (Gonen et al., 2020) is shown in Table 5 and our top-50 elements per language are shown in Table 6.

rank	De-En	De-Es	En-De	En-Es	Es-De	Es-En
1	hin	weise	each	hand	ahí	-
2	heißt	art	while	down	”	”
3	zumindest	gilt	hand	mind	-	modo
4	hohe	hinaus	upon	labour	supuesto	duda
5	dessen	sowohl	direct	while	terreno	supuesto
6	dabei	interesse	mean	due	igual	algún
7	außerdem	bleibt	secondly	form	“	bajo
8	weise	tat	light	example	incluso	pues
9	bleibt	außerdem	mind	others	luego	general
10	völlig	gesamte	beyond	another	sí	igual
11	derzeit	beim	takes	high	largo	cual
12	sorge	ebenso	kind	each	allá	luego
13	sowohl	beispiel	comes	food	origen	sí
14	darin	kommt	practice	throughout)	tipo
15	darin	letztendlich	capacity	whole	bajo	cabo
16	art	dabei	full	true	orden	alto
17	seite	davon	itself	call	precisamente	tercer
18	ihrem	sorge	open	single	corresponde	único
19	beispielsweise	heißt	sort	air	solamente	resulta
20	erster	linie	share	capacity	interés	fuera
21	macht	weder	labour	organisation	general	incluso
22	beispiel	völlig	large	power	idea	línea
23	beim	zwar	words	practice	tercer	pública
24	bzw.	erster	short	become	objeto	largo
25	ebenso	genau	individual	takes	pública	siempre
26	gilt	darin	nor	without	saben	alguna
27	bedeutet	selbst	value	nature	donde)
28	hinaus	dessen	sense	reality	alto	términos
29	voll	diejenigen	allow	comes	da	donde
30	grund	jedem	least	doubt	algún	orden
31	gesamte	macht	side	euro	cara	asimismo
32	unser	ihrem	behind	close	línea	riesgo
33	jedem	zumindest	currently	words	toda	ningún
34	steht	eindeutig	longer	far	asimismo	“
35	tatsache	sei	doubt	course	misma	mientras
36	kommt	seite	whole	circumstances	propio	vista
37	eindeutig	sorgen	down	currently	(claro
38	interesse	innerhalb	effect	non	través	efecto
39	tatsächlich	form	become	yet	fuera	práctica
40	all	handelt	outside	long	duda	régimen
41	ganze	gesamten	indeed	itself	cual	ella
42	hohen	ebene	board	least	central	misma
43	weder	jeder	free	every	alguna	público
44	aller	teil	needs	fisheries	resulta	carácter
45	gleichzeitig	beispielsweise	without	interest	claro	plan
46	gesamten	liegt	rural	general	junto	defensa
47	allein	weit	throughout	main	tipo	dentro
48	ort	schließlich	across	mean	demás	lado
49	innerhalb	übrigen	close	free	falta	número
50	jeder	reihe	another	term	siempre	personal

Table 6: Top-50 translationese words as obtained by the application of the use change concept for the three languages L1-L2 (L1 being En, Es and De) when they are translated from the other two languages L2.