

Automatic Discrimination of Human and Neural Machine Translation in Multilingual Scenarios

Malina Chichirau
Bernoulli Institute
University of Groningen
m.chichirau@student.rug.nl

Rik van Noord
CLCG
University of Groningen
rikvannoord@gmail.com

Antonio Toral
CLCG
University of Groningen
a.toral.ruiz@rug.nl

Abstract

We tackle the task of automatically discriminating between human and machine translations. As opposed to most previous work, we perform experiments in a multilingual setting, considering multiple languages and multilingual pretrained language models. We show that a classifier trained on parallel data with a single source language (in our case German–English) can still perform well on English translations that come from different source languages, even when the machine translations were produced by other systems than the one it was trained on. Additionally, we demonstrate that incorporating the source text in the input of a multilingual classifier improves (i) its accuracy and (ii) its robustness on cross-system evaluation, compared to a monolingual classifier. Furthermore, we find that using training data from multiple source languages (German, Russian, and Chinese) tends to improve the accuracy of both monolingual and multilingual classifiers. Finally, we show that bilingual classifiers and classifiers trained on multiple source languages benefit from being trained on longer text sequences, rather than on sentences.

1 Introduction

In many NLP tasks one may want to filter out machine translations (MT), but keep human translations (HT). Consider, for example, the construc-

tion of parallel corpora used for training MT systems: filtering out MT output is getting progressively harder, given the ever-increasing quality of neural MT (NMT) systems. Moreover, the existence of such high-quality NMT systems might aggravate the problem, as people are getting more likely to employ them when creating texts. In addition, it is also hard to get *fair* training data to build a classifier that can distinguish between these two types of translations, since publicly-available parallel corpora with human translations were likely used in the training of well-known publicly available MT systems (such as Google Translate or DeepL). Therefore, we believe that making the most of the scarcely available (multilingual) training data is a crucial research direction.

Previous work aiming at discriminating between HT and NMT operates mostly in a monolingual setting (Fu and Nederhof, 2021; van der Werff et al., 2022). To our knowledge, the only exception is Bhardwaj et al. (2020), who targeted English–French, and fine-tuned not only French LMs (monolingual target-only setting), but also multilingual LMs, so that the classifier had also access to the source text. However, this work used an in-house data set, therefore limiting reproducibility and practical usefulness. There is also older work that tackled statistical MT (SMT) vs HT classification (Arase and Zhou, 2013; Aharoni et al., 2014; Li et al., 2015). Nevertheless, since both the MT paradigm (SMT) and the classifiers used are not state-of-the-art anymore, less recent studies are of limited relevance today.

Compared to previous work, this paper explores the classification of HT vs NMT in the multilingual scenario in more depth, considering several languages and multilingual LMs. We demonstrate that classifiers trained on parallel data with

a single source language still work well when applied to translations from other source languages (**Experiment 1**). We show improved performance for fine-tuning multilingual LMs by incorporating the source text (**Experiment 2**), which also diminishes the gap between training and testing on different MT systems (**Experiment 3**). Moreover, we improve performance when training on additional training data from *different* source languages (**Experiment 4**) and full documents instead of isolated sentences (**Experiment 5**).

2 Method

2.1 Data

To get the source texts and human translation part of the data set, we use the data sets provided across the WMT news shared tasks of the past years.¹ As explained in the previous section, we only use the WMT test sets, to (reasonably) ensure that the popular MT systems we will be using (Google Translate and DeepL) did not use this as training data. Note that if any of the MT systems had used this data for training the task would actually be *harder*, since their translations for the data would be expected to resemble more human translations than if this data had not been used for training. An alternative would be to use in-house datasets, like Bhardwaj et al. (2020), but that also comes with an important drawback, namely limited reproducibility.

We run experiments across 7 language pairs (German, Russian, Chinese, Finnish, Gujarati, Kazakh, and Lithuanian to English) and use only the source texts that were originally written in the source language, following the findings by Zhang and Toral (2019). The data of WMT19 functions as the test set for all languages, while WMT18 is the development set (only used for German, Russian and Chinese, since the other languages are tested in a zero-shot fashion). Detailed data splits are shown in Table 1.

2.2 Translations

We obtain the MT part of the data set by translating the non-English source texts to English by using Google Translate or DeepL. The translations were obtained in November-December 2022, except for the German translations, which we take

¹For example, <https://www.statmt.org/wmt19/translation-task.html>

	Train	Dev	Test
Sentence-level			
German (WMT08-19)	8,242	1,498	2,000
Russian (WMT15-19)	4,136	1,386	1,805
Chinese (WMT17-19)	878	2,260	1,838
Finnish (WMT19)	—	—	1,996
Gujarati (WMT19)	—	—	1,016
Kazakh (WMT19)	—	—	1,000
Lithuanian (WMT19)	—	—	1,000
Document-level			
German (WMT08-19)	366	69	145
Russian (WMT15-19)	249	115	196
Chinese (WMT17-19)	123	222	163

Table 1: Number of sentences and documents per split for the languages used throughout this paper.

from van der Werff et al. (2022) and were obtained in November 2021.²

The data set we feed to our classification model is built by selecting exactly one human translation and one machine translation (either Google or DeepL) per source text. This way, we ensure there is no influence of the domain of the texts, while simultaneously ensuring a perfectly balanced data set for each experiment. Note that this also means that we actually train and test on twice as much data as is reported in Table 1. *Target-only* or *monolingual* classifiers are trained only on the English translations, while *source + target* or *multilingual* classifiers are trained on both the source text and the English translation thereof. For evaluation, we also use MT outputs from selected WMT2019’s submissions.³

2.3 Classifiers

We follow previous work (Bhardwaj et al., 2020; van der Werff et al., 2022) in fine-tuning a pre-trained language model on our task. We use DEBERTA-v3 (He et al., 2021) for the target-only classifiers since this was the best model by van der Werff et al. (2022). For the source + target classifiers we test M-BERT (Devlin et al., 2019), M-BART (Lewis et al., 2020), XLM-R (Conneau et al., 2020) and M-DEBERTA (He et al., 2021), while Bhardwaj et al. (2020) only used M-BERT and XLM-R.

²We translated the German test set in April 2023 with both Google and DeepL and compared them to the original translation of November 2021. We found BLEU scores of 98.27 and 98.54 for Google and DeepL, respectively, leading us to conclude that there are no substantial differences between the two versions of the MT systems.

³Details in Appendix A (Table 8).

↓ Eval	Trained on Google translations								Trained on DeepL translations							
	de-d	de-t	fi	gu	kk	lt	ru	zh	de-d	de-t	fi	gu	kk	lt	ru	zh
DeepL	66.0	57.4	64.8	—	—	57.6	54.6	53.8	71.7	66.9	68.7	—	—	68.6	59.5	67.7
Google	75.0	65.6	70.8	62.0	68.6	70.3	63.5	58.5	70.0	64.8	65.7	59.5	65.1	65.0	60.6	61.8
WMT ₁	57.3	70.7	67.0	65.2	66.8	62.9	58.8	58.2	60.9	66.8	66.5	60.5	64.3	66.9	58.6	65.7
WMT ₂	58.1	70.2	68.5	63.1	68.2	63.8	56.9	57.1	60.6	65.9	64.5	58.9	65.9	68.0	49.1	63.2
WMT ₃	58.9	64.9	65.2	59.4	70.9	67.0	56.4	53.7	55.7	47.1	49.2	38.6	64.5	53.9	46.2	48.5
WMT ₄	57.0	64.1	47.6	61.8	54.7	61.5	59.4	53.5	47.2	39.8	30.5	52.9	41.9	47.0	51.2	55.2

Table 2: Accuracies for the target-only DEBERTA-V3 model when training on English translations (by Google or DeepL) from German and testing on translations from a different source language and different MT system on the test set. For German we report results both on the development (de-d) and test (de-t) sets. DeepL does not offer translations from Gujarati or Kazakh.

We fine-tuned our pre-trained language models by using the Transformers library from HuggingFace (Wolf et al., 2020). We use the *ForSequence-Classification* implementations, both for the target-only as well as the source + target models. For the latter, this means that the source and target are concatenated by adding the [SEP] special character, which is the default implementation when providing two input sentences. We did experiment with adding source and target in the reverse order, but did not obtain improved performance. We did not experiment with adding a language tag to the source text.

2.4 Experimental details

The results for Experiment 1 were obtained without any hyper-parameter tuning - we simply took the settings of van der Werff et al. (2022). For finding the best multi-lingual language model (Experiment 2), we did perform a search over batch size and learning rate on the development set. We performed separate searches for the Google and DeepL translations, as well as the monolingual and bilingual settings. The final settings are shown in Table 9 in Appendix B. For Experiment 3 and Experiment 4, we used the settings of the previously found best models. For the document-level systems in Experiment 5 we used the hyperparameters listed in Table 10 in Appendix B. Reported accuracies are averaged over three runs for the sentence-level experiments (Exp 1–4) and over ten runs for the document-level experiments (Exp 5). Standard deviations (generally in range 0.2 - 2.0) are omitted for brevity, except for the document-level experiments, since they tend to be higher in the latter setting. All our code, data and results are publicly available.⁴

⁴<https://github.com/Malina03/macocu-ht-vs-mt/>

3 Results

3.1 Experiment 1: Testing on Translations from Different Source Languages

In our first experiment (with results in Table 2), we analyse the performance of our classifier when testing a target-only model on English translations from a different source language. Here, the machine translations for training our classifier come from Google or DeepL, while we evaluate on translations from Google, DeepL and the two top-ranked (WMT1, WMT2) and two bottom-ranked (WMT3, WMT4) WMT2019 submissions (Barrault et al., 2019). See Appendix A for additional details on these WMT submissions.

The results in Table 2 show that human and machine translations from a different source language can still be reasonably well distinguished. For certain languages, we are even very close to the performance on German (the original source language). The other languages do seem to show an influence of the source language, as the accuracies are generally slightly lower, but are usually still comfortably above chance-level. However, there are a few cases where the classifier now performs *below* chance level. This happened only for the bottom-ranked WMT systems (WMT3 and WMT4), which might not be representative of high-quality MT systems.

MT quality vs accuracy We are also interested in how the quality of the translations influences accuracy scores. Since we have the human (reference) translations, we plotted the accuracy score of our classifier versus an automatic MT evaluation metric, BLEU (Papineni et al., 2002), in Figure 1.⁵ What is quite striking here is that we actually obtain an increased performance for higher-quality translations. When training on DeepL translations

⁵Plots for COMET (Rei et al., 2020) instead of BLEU are in Appendix C.




Figure 1: Accuracy versus BLEU scores for each system in Table 2, using Google or DeepL translations during training.

we actually find a significant correlation between accuracy and BLEU ($R = 0.696$, $p < 0.0001$), though for Google translations we did not ($R = 0.249$, $p = 0.094$). Intuitively, it should be easier to distinguish between low-quality MT and HT, so this is likely a side-effect of training on the high-quality translations from Google and DeepL. We consider this an important lesson for future work: if a classifier learns to distinguish high-quality MT from HT, this does not mean that distinguishing lower-quality MT comes for free.

3.2 Experiment 2: Source-only vs Source+Target Classifiers

In our second experiment, we aim to determine whether having access to the source sentence improves classification performance. We test a variety of multilingual LMs, comparing their per-

	Google		DeepL	
	tgt-only	src + tgt	tgt-only	src + tgt
DEBERTA-V3	75.0	—	71.7	—
M-BERT	65.9	71.7	65.5	66.1
M-BART	69.3	71.7	61.9	68.1
XLNet-R	66.0	69.3	62.4	66.9
M-DEBERTA	70.4	74.9	65.1	71.8

Table 3: Development set accuracies of the best monolingual LM by van der Werff et al. (2022) (DEBERTA-V3) and multilingual LMs, comparing the use of target-only and source + target data. The classifiers are trained and evaluated on the German-English data (Google or DeepL). Best result per column in bold.

formance when having access only to the translation (target-only) to when also having access to the source sentence (source + target). Table 3 shows that accuracies indeed clearly improve for all of the tested LMs, with M-DEBERTA being the multilingual LM that leads to the highest accuracy. Note that this model performs similarly to the best target-only monolingual LM (DEBERTA-V3, with the scores taken from van der Werff et al. (2022)) on the development set, likely due to the higher quality of the latter LM for English. However, on the test set (also shown in Table 4), which was never seen during development of the classifiers, the multilingual model is actually clearly superior (72.3% versus 65.6%).

3.3 Experiment 3: Cross-system Evaluation

The study of van der Werff et al. (2022) showed that MT vs HT classifiers are sensitive to the MT system that was used to generate the training translations, as performance dropped considerably when doing a cross-system evaluation. However, we hypothesize that giving the classifier access to the source sentence will make it more robust to seeing translations from different MT systems at training and test times.

We show the results of the cross-MT system evaluation for the best performing target-only (DEBERTA-V3) and source + target (M-DEBERTA) models in Table 4. For training on Google and testing on DeepL, we still see a considerable drop in performance for the source + target model (around 9 points in both the dev and test sets for both

Evaluated on → ↓ Trained on	Dev		Test	
	Google	DeepL	Google	DeepL
DEBERTA-V3				
Google	75.0	66.0	65.6	57.4
DeepL	70.0	71.7	64.8	66.9
M-DEBERTA				
Google	74.9	66.2	72.3	63.8
DeepL	71.3	71.8	72.7	72.0

Table 4: Dev and test set accuracies of DEBERTA-V3 (target-only) and M-DEBERTA (source + target) when trained and evaluated on Google and DeepL. First two rows of results taken from van der Werff et al. (2022). Best score per column and classifier in bold.

the target-only and source + target classifiers). However, when training on DeepL and testing on Google, we do see a clear effect on the test set: the target-only model dropped 2.1% in accuracy (66.9 → 64.8), while the source + target model actually improved by 0.7% (72.0 → 72.7).

3.4 Experiment 4: Training on Multiple Source Languages

Here, we investigate if we can actually combine training data from different source languages to improve performance. We run experiments for German, Russian and Chinese for both the target-only and the source + target model, of which the results are shown in Table 5. Having additional training data from different source languages clearly helps, even for the multilingual source + target model. The only exception is the experiment on Chinese for the multilingual model, as the best performance (68%) is obtained by only training on the Chinese training data.⁶ There does seem to be a diminishing effect of incorporating training data from different source languages though, as the best score is only once obtained by combining all three languages as training data. Nevertheless, given the improved performance for even only small amounts of additional training data (Chinese has only 1,756 training instances), we see this as a promising direction for future work.

3.5 Experiment 5: Sentence- vs Document-level

We perform a similar experiment as van der Werff et al. (2022) by testing our classifiers on the document-level, as the WMT data sets include this

⁶The best performance on Chinese, in general, was, surprisingly, obtained by the target-only model (76.1% accuracy).

Eval → ↓ Train	DEBERTA-V3			M-DEBERTA		
	de	zh	ru	de	zh	ru
German (de)	65.6	64.2	63.3	72.3	55.1	66.1
Chinese (zh)	58.1	75.4	53.4	63.5	68.0	61.6
Russian (ru)	56.7	52.3	63.1	64.3	56.7	69.0
de + zh	66.6	76.1	63.7	72.7	66.2	68.7
de + ru	66.3	62.0	67.1	73.6	58.5	71.6
ru + zh	59.7	75.5	66.2	66.3	66.0	69.3
de + zh + ru	66.5	75.2	68.1	72.8	65.8	71.3

Table 5: Test set accuracies on discriminating between HT and Google Translate with DEBERTA-V3 (target-only) and M-DEBERTA (source + target) when training on data from one versus multiple source languages. Best score per column shown in bold.

information. We expect that the task is (a lot) easier if the classifier has access to full documents instead of just sentences. We test this with both the best monolingual (DEBERTA-V3) and multilingual (M-DEBERTA) models on Google translations from German.

Truncation DeBERTa models can in principle work with sequence lengths up to 24,528 tokens, but that does not mean this is optimal, especially when taking speed and memory requirements into account. In Table 6 we compare accuracies for different values of maximum length, or in other words, different levels of truncation. For DEBERTA-V3, the preferred truncation value is 1,024 tokens, while for M-DEBERTA we opt for 3,072. For both models, the input documents are barely truncated. The larger value for M-DEBERTA is expected, as those experiments have roughly twice the amount of input tokens (source- + target-language data versus just target data). Lengths of 3,072 (DEBERTA-V3) or 4,096 (M-DEBERTA) did not fit into our GPU memory (NVIDIA V100) even with a batch size of 1, but looking at the scores and truncation percentages, this does not seem to be an issue.

Evaluation We evaluate the models using the preferred truncation settings found above.⁷ We train on either just German, or German, Russian, and Chinese data, and evaluate on the German data.⁸ We evaluate the performance on three different classifiers: (i) applying the best sentence-level model on the documents sentence by sen-

⁷Hyperparameters used are shown in Appendix B (Table 10).

⁸Results for Russian and Chinese are in Appendix C (Table 11).

max length	DEBERTA-V3			M-DEBERTA		
	T (%)	T (avg)	Acc.	T (%)	T (avg)	Acc.
512	38	132	79.4	77	793	75.9
768	17	62	95.0	62	617	80.1
1,024	8	32	96.4	50	472	85.3
2,048	0.8	4	93.4	16	155	89.7
3,072	0.0	0.0	—	5	20	91.9

Table 6: Document-level accuracies (*Acc.*) for different values of maximum length (number of tokens) on the German development set, trained on German data. *T (%)* indicates the percentage of training documents that were truncated. *T (avg)* indicates the average amount of tokens that were truncated across the training set. Best score per classifier in bold.

tence, and taking the majority vote, (ii) simply training on the documents instead of sentences and (iii) fine-tuning the best sentence-level model on documents. The latest classifier is motivated by the fact that there are much fewer document-level training instances than there are of sentence-level (Table 1).

Document-level classifiers The results are shown in Table 7, which allows us to draw the following conclusions. For one, fine-tuning the sentence-level model on documents is clearly preferable over simply training on documents, while also comfortably outperforming the majority vote baseline. Fine-tuning not only leads to the highest accuracies, but also to the lowest standard deviations, indicating that this classifier is more stable than the other two. Second, we confirm our two previous findings: the models can improve performance when training on texts from a different source language (Chinese and Russian in this case) and the models clearly benefit from having access to the source text itself during training and evaluation.

4 Conclusion

This paper has investigated the discrimination between neural machine translation (NMT) and human translation (HT) in multilingual scenarios, using as classifiers monolingual and multilingual language models (LMs) that are fine-tuned with small amounts of task-specific labelled data.

We have found out that a monolingual classifier trained on English translations from a given source language still performs well above chance on English translations from other source languages. Using a multilingual LM and therefore having access also to the source sentence results overall in better performance than an equivalent LM that only

Trained on →	German (de)		de + ru + zh	
	DEB	M-DEB	DEB	M-DEB
Majority vote	68.5 ±8.7	73.1 ±5.7	75.6 ±6.7	76.5 ±4.7
Doc-level	62.6 ±3.6	75.3 ±3.9	67.3 ±10.7	83.0 ±2.2
Doc-level (ft)	81.1 ±2.7	86.0 ±1.2	87.0 ±2.6	88.7 ±1.4

Table 7: Document-level accuracies and standard deviations with DEBERTA-V3 (target-only, denoted as DEB) and M-DEBERTA (source + target, M-DEB) when evaluating on the test that has German as the source language using Google as the MT system. Best result per column shown in bold.

has access to the target sentence. Such a classifier seems more robust in a cross-system situation, i.e. when the MT systems used to train and evaluate the classifier are different. Moreover, as task-specific data is limited, we experimented with (i) training on data from different source languages and (ii) training on the document-level instead of the sentence-level, with improved performance in both settings.

4.1 Future work

In this work, we took an important step toward developing an accurate, reliable, and accessible classifier that can distinguish between HT and MT. There are, of course, still many research directions to explore, in particular regarding combining different source languages and MT systems during training. Moreover, in many practical applications, it is unknown whether a text is actually a translation, as it can also be an original text. Therefore, in future work, we aim to develop a classifier that can distinguish between original texts, human translations, and machine translations.

Acknowledgements

The authors received funding from the European Union’s Connecting Europe Facility 2014-2020 - CEF Telecom, under Grant Agreement No. INEA/CEF/ICT/A2020/2278341 (MaCoCu). This communication reflects only the authors’ views. The Agency is not responsible for any use that may be made of the information it contains. We thank the Center for Information Technology of the University of Groningen for providing access to the Peregrine high performance computing cluster. Finally, we thank all our MaCoCu colleagues for their valuable feedback throughout the project.

References

- Aharoni, Roei, Moshe Koppel, and Yoav Goldberg. 2014. Automatic detection of machine translated text and translation quality estimation. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 289–295.
- Arase, Yuki and Ming Zhou. 2013. Machine translation detection from monolingual web-text. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1597–1607.
- Barrault, Loïc, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, Mathias Müller, Santanu Pal, Matt Post, and Marcos Zampieri. 2019. Findings of the 2019 conference on machine translation (WMT19). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, Florence, Italy, August. Association for Computational Linguistics.
- Bawden, Rachel, Nikolay Bogoychev, Ulrich Germann, Roman Grundkiewicz, Faheem Kirefu, Antonio Valerio Miceli Barone, and Alexandra Birch. 2019. The University of Edinburgh’s submissions to the WMT19 news translation task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 103–115, Florence, Italy, August. Association for Computational Linguistics.
- Bei, Chao, Hao Zong, Conghu Yuan, Qingming Liu, and Baoyong Fan. 2019. GTCOM neural machine translation systems for WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 116–121, Florence, Italy, August. Association for Computational Linguistics.
- Bhardwaj, Shivendra, David Alfonso Hermelo, Phillippe Langlais, Gabriel Bernier-Colborne, Cyril Goutte, and Michel Simard. 2020. Human or neural translation? In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6553–6564, Barcelona, Spain (Online), December. International Committee on Computational Linguistics.
- Biçici, Ergun. 2019. Machine translation with parfda, Moses, kenlm, nplm, and PRO. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 122–128, Florence, Italy, August. Association for Computational Linguistics.
- Briakou, Eleftheria and Marine Carpuat. 2019. The University of Maryland’s Kazakh-English neural machine translation system at WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 134–140, Florence, Italy, August. Association for Computational Linguistics.
- Budiwati, Sari Dewi, Al Hafiz Akbar Maulana Siagian, Tirana Noor Fatyanosa, and Masayoshi Aritsugi. 2019. DBMS-KU interpolation for WMT19 news translation task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 141–146, Florence, Italy, August. Association for Computational Linguistics.
- Conneau, Alexis, 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, July. Association for Computational Linguistics.
- Dabre, Raj, Kehai Chen, Benjamin Marie, Rui Wang, Atsushi Fujita, Masao Utiyama, and Eiichiro Sumita. 2019. NICT’s supervised neural machine translation systems for the WMT19 news translation task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 168–174, Florence, Italy, August. Association for Computational Linguistics.
- Devlin, Jacob, 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, June. Association for Computational Linguistics.
- Fu, Yingxue and Mark-Jan Nederhof. 2021. Automatic classification of human translation and machine translation: A study from the perspective of lexical diversity. In *Proceedings for the First Workshop on Modelling Translation: Translatology in the Digital Age*, pages 91–99, online, May. Association for Computational Linguistics.
- Goyal, Vikrant and Dipti Misra Sharma. 2019. The IIIT-H Gujarati-English machine translation system for WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 191–195, Florence, Italy, August. Association for Computational Linguistics.
- Guo, Xinze, Chang Liu, Xiaolong Li, Yiran Wang, Guoliang Li, Feng Wang, Zhitao Xu, Liuyi Yang, Li Ma, and Changliang Li. 2019. Kingsoft’s neural machine translation system for WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 196–202, Florence, Italy, August. Association for Computational Linguistics.
- He, Pengcheng, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style

- pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.
- Lewis, Mike, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online, July. Association for Computational Linguistics.
- Li, Zhenhao and Lucia Specia. 2019. A comparison on fine-grained pre-trained embeddings for the WMT19Chinese-English news translation task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 249–256, Florence, Italy, August. Association for Computational Linguistics.
- Li, Yitong, Rui Wang, and Hai Zhao. 2015. A machine learning method to distinguish machine translation from human translation. In *Proceedings of the 29th Pacific Asia Conference on Language, Information and Computation: Posters*, pages 354–360.
- Li, Bei, Yinqiao Li, Chen Xu, Ye Lin, Jiqiang Liu, Hui Liu, Ziyang Wang, Yuhao Zhang, Nuo Xu, Zeyang Wang, Kai Feng, Hexuan Chen, Tengbo Liu, Yanyang Li, Qiang Wang, Tong Xiao, and Jingbo Zhu. 2019. The NiuTrans machine translation systems for WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 257–266, Florence, Italy, August. Association for Computational Linguistics.
- Molchanov, Alexander. 2019. PROMT systems for WMT 2019 shared translation task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 302–307, Florence, Italy, August. Association for Computational Linguistics.
- Mondal, Riktim, Shankha Raj Nayek, Aditya Chowdhury, Santanu Pal, Sudip Kumar Naskar, and Josef van Genabith. 2019. JU-Saarland submission to the WMT2019 English–Gujarati translation shared task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 308–313, Florence, Italy, August. Association for Computational Linguistics.
- Ng, Nathan, 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, August. Association for Computational Linguistics.
- Papineni, Kishore, 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 of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA, July. Association for Computational Linguistics.
- Pinnis, Marcis, Rihards Krišlauks, and Matīss Rikters. 2019. Tilde’s machine translation systems for WMT 2019. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 327–334, Florence, Italy, August. Association for Computational Linguistics.
- Pirinen, Tommi. 2019. Apertium-fin-eng–rule-based shallow machine translation for WMT 2019 shared task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 335–341, Florence, Italy, August. Association for Computational Linguistics.
- Rei, Ricardo, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online, November. Association for Computational Linguistics.
- Rosendahl, Jan, Christian Herold, Yunsu Kim, Miguel Graça, Weiyue Wang, Parnia Bahar, Yingbo Gao, and Hermann Ney. 2019. The RWTH Aachen University machine translation systems for WMT 2019. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 349–355, Florence, Italy, August. Association for Computational Linguistics.
- Sun, Meng, Bojian Jiang, Hao Xiong, Zhongjun He, Hua Wu, and Haifeng Wang. 2019. Baidu neural machine translation systems for WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 374–381, Florence, Italy, August. Association for Computational Linguistics.
- van der Werff, Tobias, Rik van Noord, and Antonio Toral. 2022. Automatic discrimination of human and neural machine translation: A study with multiple pre-trained models and longer context. In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*, pages 161–170, Ghent, Belgium, June. European Association for Machine Translation.
- Wolf, Thomas, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, October. Association for Computational Linguistics.
- Xia, Yingce, Xu Tan, Fei Tian, Fei Gao, Di He, Weicong Chen, Yang Fan, Linyuan Gong, Yichong

Leng, Renqian Luo, Yiren Wang, Lijun Wu, Jinhua Zhu, Tao Qin, and Tie-Yan Liu. 2019. Microsoft Research Asia’s systems for WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 424–433, Florence, Italy, August. Association for Computational Linguistics.

Zhang, Mike 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, August. Association for Computational Linguistics.

A WMT MT Systems

Table 8 shows the specific WMT19 systems that were used during Experiment 1. Barrault et al. (2019) did not specify which specific online systems were used.

B Hyperparameters

Sentence-level hyperparameters used in our experiments are shown in Table 9, while the document-level settings are shown in Table 10.

C Additional Results

Figure 2 shows additional results for Experiment 1, specifically scatter plots of the accuracy of the classifier versus COMET scores for each system for both Google and DeepL. This complements Figure 1 in Section 3.1, in which BLEU was used instead of COMET. The trends are very similar in both figures.

Table 11 shows additional evaluation results on document-level classification (Experiment 5), as opposed to Table 7 in which we evaluated on the test set that has German as the source language. We observe that fine-tuning the sentence-level model is still generally preferable, though there are few cases in which just training on documents resulted in the best performance. A curious observation is that for Chinese including the source text does generally not lead to improved performance, while this is not the case for Russian and German.

	WMT1	WMT2	WMT3	WMT4
de	Ng et al. (2019)	Rosendahl et al. (2019)	Molchanov (2019)	online-X
fi	Xia et al. (2019)	online-Y	Biçici (2019)	Pirinen (2019)
gu	Li et al. (2019)	Bawden et al. (2019)	Goyal and Sharma (2019)	Mondal et al. (2019)
kk	online-B	Li et al. (2019)	Briakou and Carpuat (2019)	Budiwati et al. (2019)
lt	Bei et al. (2019)	Pinnis et al. (2019)	JUMT	online-X
ru	Ng et al. (2019)	online-G	online-X	Dabre et al. (2019)
zh	Sun et al. (2019)	Guo et al. (2019)	Li and Specia (2019)	online-X

Table 8: WMT systems used in our analysis. WMT1 and WMT2 are the two top-ranked systems, while WMT3 and WMT4 are the two bottom-ranked systems. The JUMT system did not submit a paper.

	Monolingual				Multilingual			
	Learning Rate		Batch Size		Learning Rate		Batch Size	
	Google	DeepL	Google	DeepL	Google	DeepL	Google	DeepL
DEBERTA-V3	$1e^{-5}$	$1e^{-5}$	32	32	—	—	—	—
M-BERT	$1e^{-5}$	$1e^{-5}$	16	32	$1e^{-5}$	$1e^{-5}$	16	32
M-BART	$1e^{-5}$	$5e^{-6}$	16	32	$5e^{-6}$	$1e^{-5}$	16	16
XLM-R	$1e^{-5}$	$1e^{-5}$	16	32	$1e^{-5}$	$1e^{-5}$	16	16
M-DEBERTA	$1e^{-5}$	$1e^{-5}$	32	32	$5e^{-5}$	$1e^{-5}$	32	16

Table 9: Final hyper-parameter settings for the models used throughout the paper. We experimented with a batch size of {16, 32, 64} and a learning rate of $\{1e^{-4}, 1e^{-5}, 5e^{-5}, 1e^{-6}, 5e^{-6}\}$.

	Max Sequence Length	Learning Rate	Batch Size	Gradient Accumulation
DEBERTA-V3	1,024	$1e^{-5}$	2	8
M-DEBERTA	3,072	$1e^{-5}$	1	8

Table 10: Final hyper-parameter settings for the models trained on the document level. We experimented with batch sizes of $\{1, 2, 4, 8\}$ and different gradient accumulation values such that the effective batch size was at most 16 due to the hardware limitations. The learning rates tested were $\{1e^{-6}, 2e^{-6}, 5e^{-6}, 1e^{-5}\}$.

Tested on →	Russian				Chinese			
	German (de)		de + ru + zh		German (de)		de + ru + zh	
	DEB	M-DEB	DEB	M-DEB	DEB	M-DEB	DEB	M-DEB
Majority vote	67.2 ±9.1	68.8 ±2.2	71.0 ±5.5	69.5 ±4.7	62.5 ±10.8	53.8 ±5.2	92.6 ±11.9	65.2 ±4.6
Doc-level	58.7 ±3.8	73.3 ±2.7	66.0 ±6.4	71.7 ±2.4	53.6 ±3.6	69.0 ±8.8	84.8 ±13.6	84.4 ±6.9
Doc-level (ft)	78.4 ±1.8	72.8 ±2.4	75.3 ±2.3	80.6 ±1.5	76.5 ±6.4	52.6 ±1.2	96.0 ±0.9	78.0 ±1.8

Table 11: Document-level accuracies when evaluating on test sets where the source language is either Russian or Chinese (see Table 7 for results on German). We train source-only DEBERTA-V3 (DEB) and source + target M-DEBERTA (M-DEB) models on either German, or German, Russian and Chinese combined.




Figure 2: Accuracy versus COMET scores for each system in Table 2, using Google (left) or DeepL (right) translations during training. Accuracy versus BLEU scores can be found in Figure 1.