

Exploring Diversity in Back Translation for Low-Resource Machine Translation

Laurie Burchell and Alexandra Birch and Kenneth Heafield

Institute for Language, Cognition, and Computation

School of Informatics, University of Edinburgh

10 Crichton Street, Edinburgh, EH8 9AB, UK

{laurie.burchell, a.birch, kenneth.heafield}@ed.ac.uk

Abstract

Back translation is one of the most widely used methods for improving the performance of neural machine translation systems. Recent research has sought to enhance the effectiveness of this method by increasing the ‘diversity’ of the generated translations. We argue that the definitions and metrics used to quantify ‘diversity’ in previous work have been insufficient. This work puts forward a more nuanced framework for understanding diversity in training data, splitting it into lexical diversity and syntactic diversity. We present novel metrics for measuring these different aspects of diversity and carry out empirical analysis into the effect of these types of diversity on final neural machine translation model performance for low-resource English↔Turkish and mid-resource English↔Icelandic. Our findings show that generating back translation using nucleus sampling results in higher final model performance, and that this method of generation has high levels of both lexical and syntactic diversity. We also find evidence that lexical diversity is more important than syntactic for back translation performance.

1 Introduction

The data augmentation technique of back translation (BT) is used in nearly every current neural machine translation (NMT) system to reach optimal performance (Edunov et al., 2020; Barrault et al., 2020; Akhbardeh et al., 2021, *inter alia*). It involves creating a pseudo-parallel dataset by translating target-side monolingual data into the source language using a secondary NMT system (Sennrich et al., 2016). In this way, it enables the incorporation of monolingual data into the NMT system. Whilst adding data in this way helps nearly all language pairs, it is particularly important for low-resource NMT where parallel data is scarce by definition.

Because of its ubiquity, there has been extensive research into how to improve BT (Burlot and Yvon,

2018; Hoang et al., 2018; Fadaee and Monz, 2018; Caswell et al., 2019), especially in ways which increase the ‘diversity’ of the back-translated dataset (Edunov et al., 2018; Soto et al., 2020). Previous work (Gimpel et al., 2013; Ott et al., 2018; Vanmassenhove et al., 2019) has found that machine translations lack the diversity of human productions. This is because most translation systems use some form of maximum a-posteriori (MAP) estimation, meaning that they will always favour the most probable output. Edunov et al. (2018) and Soto et al. (2020) argue that this makes standard BT data worse training data since it lacks ‘richness’ or diversity.

Despite the focus on increasing diversity in BT, what ‘diversity’ actually means in the context of NMT training data is ill-defined. In fact, Tevet and Berant (2021) point out that there is no standard metric for measuring diversity. Most previous work uses the BLEU score between candidate sentences or another n-gram based metric to estimate similarity (Zhu et al., 2018; Hu et al., 2019; He et al., 2018; Shen et al., 2019; Shu et al., 2019; Holtzman et al., 2020; Thompson and Post, 2020). However, such metrics mostly measure changes in the vocabulary or spelling. Because of this, they are likely to be less sensitive to other kinds of variety such as changes in structure.

We argue that quantifying ‘diversity’ using n-gram based metrics alone is insufficient. Instead, we split diversity into two aspects: variety in the word choice and spelling, and variety in structure. We call these aspects *lexical diversity* and *syntactic diversity* respectively. Here, we follow recent work in natural language generation and particularly paraphrasing (e.g. Iyyer et al., 2018; Krishna et al., 2020; Goyal and Durrett, 2020; Huang and Chang, 2021; Hosking and Lapata, 2021) which explicitly models the meaning and form of the input separately. Of course, there are likely more kinds of diversity than this, but this distinction provides

a common-sense framework to extend our understanding of the concept. To our knowledge, no other previous work in data augmentation has attempted to isolate and automatically measure syntactic and lexical diversity.

Building from our definition, we introduce novel metrics aimed at measuring lexical and syntactic diversity separately. We then carry out an empirical study into what effect training data with these two kinds of diversity has on final NMT performance in the context of low-resource machine translation. We do this by creating BT datasets using different generation methods and measuring their diversity. We then evaluate what impact different aspects of diversity have on final model performance. We find that a high level of diversity is beneficial for final NMT performance, though lexical diversity seems more important than syntactic diversity. Importantly though there are limits to both; the data should not be so ‘diverse’ that it affects the adequacy of the parallel data.

We summarise our contributions as follows:

- We put forward a more nuanced definition of ‘diversity’ in NMT training data, splitting it into *lexical diversity* and *syntactic diversity*. We present two novel metrics for measuring these different aspects of diversity.
- We carry out empirical analysis into the effect of these types of diversity on final NMT model performance for low-resource English↔Turkish and mid-resource English↔Icelandic.
- We find that nucleus sampling is the highest-performing method of generating BT, and it combines both lexical and syntactic diversity.
- We make our code publicly available.¹

2 Methods

We explain each method we use for creating diverse BT datasets in Section 2.1, then discuss our metrics for diversity in Section 2.2.

2.1 Generating diverse back translation

We use four methods to generate diverse BT datasets: beam search, pure sampling, nucleus sampling, and syntax-group fine-tuning. The first three were chosen because they are in common

use and so more relevant for future work. The last, syntax-group fine-tuning, aims to increase syntactic diversity specifically and so allows us to separate its effect on final NMT performance from lexical diversity. For each method, we create a diverse BT dataset by generating three candidate translations for each input sentence. This allows us to measure diversity whilst keeping the ‘meaning’ of the sentence as similar as possible. In this way, we measure inter-sentence diversity as a proxy for the diversity of the dataset as a whole. We discuss our datasets in detail in Section 3.1.

Beam search Beam search is the most common search algorithm used to decode in NMT systems. Whilst it is generally successful in finding a high-probability output, the translations it produces tend to lack diversity since it will always default to the most likely alternative in the case of ambiguity (Ott et al., 2018). We use beam search to generate three datasets for each language pair, using a beam size of five and no length penalty:

- *base*: three million input sentences used to generate one output per input (BT dataset length: three million)
- *beam*: three million input sentences used to generate three outputs per input (BT dataset length: nine million)
- *base-big*: nine million input sentences used to generate one output per output (BT dataset length: nine million)

Pure sampling An alternative to beam search is sampling from the model distribution. At each decoding step, we sample from the learned distribution without restriction to generate output. This method means we are likely to generate a much wider range of tokens than restricting our choice to those which are most likely (as in beam search). However, it also means that the generated text is less likely to be adequate (have the same meaning as the input) as the output space does not necessarily restrict itself to choices which best reflect the meaning of the input. In other words, the output may be diverse, but it may not be the kind of diversity that we want for NMT training data.

We create one dataset per language pair (*sampling*) by generating three candidate translations for each of the three million monolingual input sentences. This results in nine-million line

¹github.com/laurieburchell/exploring-diversity-bt

BT dataset. We set our beam size to five when generating.

Nucleus sampling Nucleus or top-p sampling is another sampling-based method, introduced by Holtzman et al. (2020). Unlike pure sampling, which samples from the entire distribution, top-p sampling only samples from the highest probability tokens whose cumulative probability mass exceeds the pre-chosen threshold p . The intuition is that when only a small number of tokens are likely, we want to limit our sampling space to those. However, when there are many likely hypotheses, we want to widen the number of tokens we might sample from. We chose this method in the hope it represents a middle ground between high-probability but repetitive beam search generations, and more diverse but potentially low-adequacy pure sampling generation. We create one dataset per language pair (*nucleus*) by generating three hypothesis translations for each of the three million monolingual input sentences. Each dataset is therefore nine million lines long. We set the beam size to five and p to 0.95.

Syntax-group fine-tuning For our analysis in this paper, we want to generate diverse BT in a way which focuses on syntactic diversity over lexical diversity, so that we can separate out its effect on final NMT performance. We therefore take a fine-tuning approach for our final generation method. To do this, we generate the dependency parse of each sentence in the English side of the parallel data for each language pair using the Stanford neural network dependency parser (Chen and Manning, 2014). We then label each pair of parallel sentences in the training data according to the first split in the corresponding syntactic parse tree. We then create three fine-tuning training datasets out of the three biggest syntactic groups.² Finally, we take NMT models trained on parallel data alone and restart training on each syntactic-group dataset, resulting in three NMT systems which are fine-tuned to produce a particular syntactic structure. We are only able to create models this way which translate into English, as good syntactic parsers are not available for the other languages in our study.

To verify this method works as expected, we translated the test set for each language pair with the model trained on parallel data only. We then

²For English–Turkish, we combine the third and fourth largest syntactic groups to create the third fine-tuning dataset, as the third-largest syntactic group alone was not large enough for successful fine-tuning.

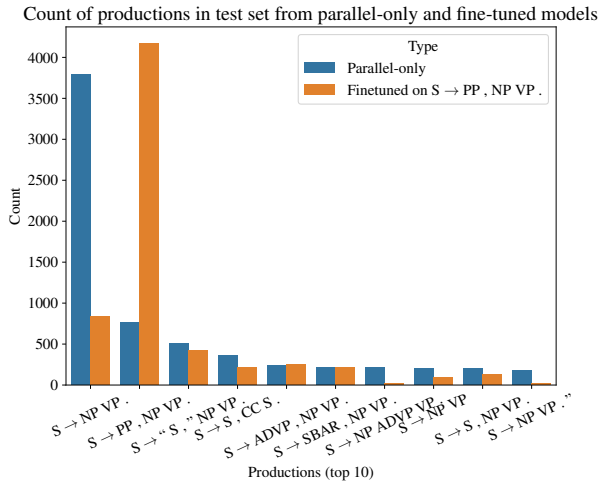


Figure 1: The count of the top-ten syntactic groups produced by the parallel-only Turkish→English NMT model compared to the number of those productions produced by a Turkish→English NMT model fine-tuned on the second-most common syntactic group ($S \rightarrow PP NP VP .$). The fine-tuned model produces more examples of the required syntactic group. Input data is the combined WMT test sets.

translated the same test set with each fine-tuned model and checked it was producing more of the required syntactic group. We did indeed find that fine-tuning resulted in more candidate sentences from the required group. Figure 1 gives an example of the different pattern of productions between the parallel-only model and a model fine-tuned on a particular syntactic group ($S \rightarrow PP NP VP .$)

2.2 Diversity metrics

We use three primary metrics to measure lexical and syntactic diversity: i-BLEU, i-chrF, and tree kernel difference. As mentioned in Section 2.1, we generate three output sentences for each input to our BT systems and measure inter-sentence diversity as a proxy for the diversity produced by the system. Due to compute time, we calculate all inter-sentence metrics over a sample of 30,000 sentence groups rather than the whole BT dataset.

i-BLEU Following previous work, we calculate the BLEU score between all sentence pairs generated from the same input (Papineni et al., 2002), take the mean and then subtract it from one to give inter-sentence or i-BLEU (Zhu et al., 2018). We believe that lexical diversity as we define it is the main driver of this metric, since BLEU scores are calculated based on n-gram overlap and so the biggest changes to the score will result from changes to the

words used (though changes in ordering of words and their morphology will also have an effect). The higher the i-BLEU score, the higher the diversity of output.

i-chrF Building from i-BLEU, we introduce i-chrF, which is generated in the same way as i-BLEU but using chrF (Popović, 2015). Since chrF is also based on n-gram overlap, we believe it will also mostly measure lexical diversity. However, i-chrF is based on character rather than word overlap, and so should be less affected by morphological changes to the form of words than i-BLEU. We calculate both chrF and BLEU scores using the sacreBLEU toolkit (Post, 2018).

Tree kernel difference We propose a novel metric which focuses on syntactic diversity: mean tree kernel difference. To calculate it, we first generate the dependency parse of each candidate sentence using the Stanford neural network dependency parser (Chen and Manning, 2014). We replace all terminals with a dummy token to minimise the effect of lexical differences, then we calculate the tree kernel for each pair of parses using code from Conklin et al. (2021), which is in turn based on Moschitti (2006). Finally, we calculate the mean across all pairs to give the mean tree kernel difference for each set of generated sentences.

We are only able to calculate the tree kernel metric for the English datasets due to the lack of reliable parsers in Turkish and Icelandic, though this method could extend to any language with a reasonable parser available. The higher the score, the higher the diversity of the output.

Summary statistics We calculate mean word length, mean sentence length, and vocabulary size over the entire generated dataset as summary statistics. We use the definition of ‘word’ as understood by the bash `wc` command to calculate all metrics, since we are only interested in a rough measure to check for degenerate results.

3 Experiments

Having discussed the methods by which we generate diverse BT datasets and the metrics with which we measure the diversity in these datasets, we now outline our experimental set up for testing the effect of training data diversity on final NMT model performance.

3.1 Data and preprocessing

We carry out our experiments on two language pairs: low-resource Turkish–English and mid-resource Icelandic–English. These languages are sufficiently low-resource that augmenting the training data will likely be beneficial, but well-resourced enough that we can still train a reasonable back-translation model on the available parallel data alone.

Data provenance The Turkish–English parallel data is from the WMT 2018 news translation task (Bojar et al., 2018). The training data is from the SETIMES dataset, a parallel dataset of news articles in Balkan languages (Tiedemann, 2012). We use the development set from WMT 2016 and the test sets from WMT 2016–18.

The Icelandic–English parallel data is from the WMT 2021 news translation task (Akhbardeh et al., 2021). There are four sources of training data: ParIce (Barkarson and Steingrímsson, 2019), filtered as described in Jónsson et al. (2020); Paracrawl (Bañón et al., 2020); WikiMatrix (Schwenk et al., 2021); and WikiTitles³. We use the development and test sets provided for WMT 2021.

The English monolingual data is made up of news crawl data from 2016 to 2020, version 16 of news-commentary crawl,⁴ and crawled news discussions from 2012 to 2019.⁵ The Turkish monolingual data is news crawl data from 2016 to 2020.⁶ The Icelandic monolingual data is made up of news crawl data from 2020, and part one of the Icelandic Gigaword dataset (Steingrímsson et al., 2018).

Data cleaning Our cleaning scripts are adapted from those provided by the Bergamot project.⁷ The full data preparation procedure is provided in the repo accompanying this paper. After cleaning, the Turkish–English parallel dataset contains 202 thousand lines and the Icelandic–English parallel dataset contains 3.97 million lines. The English, Icelandic, and Turkish cleaned monolingual datasets contain 487 million, 39.9 million, and 26.1 million lines respectively. We select 9 million lines of each monolingual dataset for BT at random since all the monolingual datasets are the same domain as the test sets.

³data.statmt.org/wikititles/v3

⁴data.statmt.org/news-commentary/v16

⁵data.statmt.org/news-discussions/en

⁶data.statmt.org/news-crawl

⁷github.com/browsermt/students/tree/master/train-student

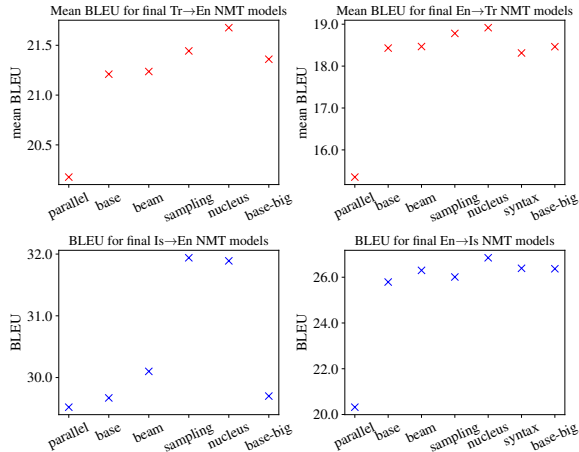


Figure 2: Mean BLEU score on WMT test sets for English↔Turkish and English↔Icelandic models trained on different BT datasets. For English↔Turkish, we give the mean score on WMT 16, WMT 17, and WMT 18 test sets. For English↔Icelandic, we give the score on the WMT 21 test set.

Text pre-processing We learn a joint BPE model with SentencePiece using the concatenated training data for each language pair (Kudo and Richardson, 2018). We set vocabulary size to 16,000 and character coverage to 1.0. All other settings are default. We apply this model to the training, development, and test data. We remove the BPE segmentation before calculating any metrics.

3.2 Model training

Model architecture and infrastructure All NMT models in this paper are transformer models (Vaswani et al., 2017). We give full details about hyper-parameters and infrastructure in Appendix A.2.

Parallel-only models for back translation For each language pair and in both directions, we train an NMT model on the cleaned parallel data alone using the relevant hyper-parameter settings in Table 5. We measure the performance of these models by calculating the BLEU score (Papineni et al., 2002) using the sacreBLEU toolkit (Post, 2018)⁸ and by evaluating the translations with COMET using the wmt20-comet-da model (Rei et al., 2020).

Generating back translation For each language pair and in each direction, we use the trained parallel-only models to generate back translation

⁸BLEU|nrefs:1|case:mixed|eff:no|tok:13a|smooth:exp|version:2.0.0

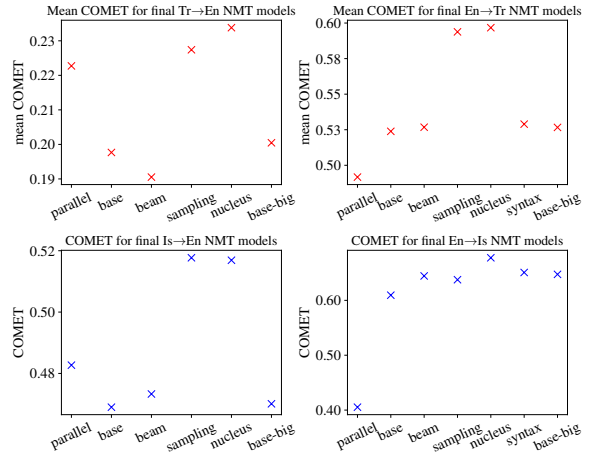


Figure 3: Mean COMET score on WMT test sets for English↔Turkish and English↔Icelandic models trained on different BT datasets. For English↔Turkish, we give the mean score on WMT 16, WMT 17, and WMT 18 test sets. For English↔Icelandic, we give the score on the WMT 21 test set.

datasets as described in Section 2.1. We translate the same three million sentences of monolingual data each time for consistency, translating an additional six million lines of monolingual data for the *base-big* dataset.

Training final models We train final models for each language direction on the concatenation of the parallel data and each back-translation dataset (back-translation on the source side, original monolingual data as target). We measure the final performance of these models using BLEU and COMET as before.

4 Results and Analysis

4.1 Final model performance

Figures 2 and 3 show the mean BLEU and COMET scores achieved by the final models trained on the concatenation of the parallel data and the different BT datasets. In most cases, adding any BT data to the training data results in some improvement over the parallel-only baseline for both scores. However, augmenting the training data with BT produced with nucleus sampling nearly always results in the strongest performance, with mean gains of 2.88 BLEU or 0.078 COMET. This compares to mean gains of 2.24 BLEU or 0.026 COMET when using the baseline BT dataset of three million lines translated with beam search. Pure sampling tends to perform similarly but not quite as well as nucleus sampling. Based on this result, we suggest that

Dataset	<i>base-big</i>	<i>beam</i>	<i>sampling</i>	<i>nucleus</i>
Sent. len.	12.23	12.21	13.11	12.74
Word. len.	8.19	8.17	8.37	8.28
Vocab	1.6M	1.0M	5.6M	3.4M
i-BLEU	-	38.11	86.69	83.27
i-chrF	-	17.91	58.84	53.95

Table 1: Diversity metrics for the Turkish BT datasets (original language: English) used to train the Tr→En models. Inter-sentence metrics are calculated on a sample of 30k triplets. ‘M’ = million.

Dataset	<i>base-big</i>	<i>beam</i>	<i>sampling</i>	<i>nucleus</i>
Sent. len.	15.65	14.79	14.92	14.73
Word len.	6.54	6.91	7.33	7.15
Vocab.	1.3M	0.82M	11M	5.6M
i-BLEU	-	30.89	86.41	79.67
i-chrF	-	16.09	66.06	57.83

Table 2: Diversity metrics for the Icelandic BT datasets (original language: English) used to train the Is→En models. Inter-sentence metrics are calculated on a sample of 30k triplets. ‘M’ = million.

future work generate BT with nucleus sampling rather than pure sampling.

4.2 Diversity metrics

We give the diversity metrics for each language pair and each generated dataset in Tables 1 to 4.⁹ Sentence and word lengths are comparable across the same language for all generation methods, suggesting that each method is generating tokens from roughly the right language distribution. However, the vocabulary size is much larger for *nucleus* compared to *base* or *beam*, and *sampling* is around twice that of *nucleus*. Examining the data, we find many neologisms (that is, ‘words’ which do not appear in the training data) for *nucleus* and more still for *sampling*. We note that the *syntax-groups* dataset has a much smaller vocabulary again; this is what we would hope if the generation method is producing syntactic rather than lexical diversity as required. We give representative examples of generated triples in Appendix A.1, along with some explanation of how the phenomena they demonstrate fit into the general trend of the dataset.

Effect on performance With respect to the inter-sentence diversity metrics (i-BLEU, i-chrF, and tree kernel scores), we see that the *sampling* dataset has the highest diversity scores, followed by *nucleus*,

⁹We omit *base* for reasons of space and because its different length to the other datasets makes comparison difficult (3 million lines compared to 9 million for the others).

Dataset	<i>base+</i>	<i>beam</i>	<i>sampl.</i>	<i>nucleus</i>	<i>syntax</i>
Sent. len.	16.98	17.03	17.85	17.54	17.28
Word len.	6.05	6.04	6.25	6.11	6.06
Vocab	0.89M	0.54M	4.9M	2.5M	0.64M
i-BLEU	-	30.74	83.52	78.92	42.26
i-chrF	-	16.28	57.16	51.82	23.32
Kernel	-	72.20	97.33	95.91	83.43

Table 3: Diversity metrics for the English BT datasets (original language: Turkish) used to train the En→Tr models. Inter-sentence metrics are calculated on a sample of 30k triplets. ‘M’ = million.

Dataset	<i>base+</i>	<i>beam</i>	<i>sampl.</i>	<i>nucleus</i>	<i>syntax</i>
Sent. len.	20.45	22.75	21.34	21.13	18.29
Word len.	5.83	5.83	6.33	6.08	5.89
Vocab.	0.66M	0.41M	12M	5.6M	0.49M
i-BLEU	-	22.75	92.31	88.86	77.17
i-chrF	-	11.95	72.20	67.16	56.90
Kernel	-	65.72	99.35	98.74	99.40

Table 4: Diversity metrics for the English BT datasets (original language: Icelandic) used to train the En→Is models. Inter-sentence metrics are calculated on a sample of 30k triplets. ‘M’ = million.

then *syntax*, then *beam*. Taken together with the performance scores and the summary statistics, this suggests that NMT data benefits from a high level of diversity, but not so high that the two halves of the parallel data no longer have the same meaning (as shown by the very high vocabulary size for *sampling*).

Metric correlation There is a high correlation between i-BLEU, i-chrF, and tree kernel score for the *beam*, *sampling*, and *nucleus* datasets. This is not entirely unexpected: it is likely to be difficult if not impossible to disentangle lexical and syntactic diversity, since changing sentence structure would also affect the word choice and vice versa.

This correlation is much weaker for the *syntax-groups* dataset: whilst the tree-kernel scores are comparable to the *sampling* and *nucleus* datasets, there is a much smaller increase in the other (lexical) diversity scores. This suggests that this generation method encourages relatively more syntactic variation than lexical compared to the other diverse generation method, as was its original aim (see paragraph on syntax-group fine-tuning in section 2.1). The fact that the final model trained on this BT dataset has lower performance compared to other forms of diversity suggests that lexical diversity is more important than syntactic diversity when undertaking data augmentation. We leave it

to future work to investigate this hypothesis further.

4.3 Data augmentation versus more monolingual data

The right-most cross in each quadrant of Figures 2 and 3 gives the performance of *base-big*, the dataset where we simply add six million more lines of new data rather than carrying out data augmentation. Interestingly, pure and nucleus sampling both often outperform *base-big*. This may be because the model over-fits to too much back-translated data, whereas having multiple sufficiently-diverse pseudo-source sentences for each target sentence has a regularising effect on the model.

To further support this hypothesis, Figure 4 gives training perplexity for the first 50,000 steps of training for the final Icelandic→English models, which are representative of the results for the other language pairs. We see that the *base-big* dataset has the lowest training perplexity at each step, suggesting this data is easier to model. Conversely, the model has highest training perplexity on the *sampling* and *nucleus* datasets, suggesting generating the data this way has a regularising effect.




Figure 4: Mean training perplexity for the first 50 thousand steps of training for final English→Turkish models. The model has highest training perplexity on the *sampling* then *nucleus* datasets. The lowest training perplexity is on the *beam* and *base-big* datasets.

4.4 Translationese effect

Several studies have found that back-translated text is easier to translate than forward-translated text, and so inflates intrinsic metrics like BLEU (Edunov et al., 2020; Graham et al., 2020; Roberts et al., 2020). To use a concrete example, the WMT test sets for English to Turkish are made up of half native English translated into Turkish, and half native

Turkish translated into English. We want models that perform well when translating *from* native text (in this example: the native English side), as this is the usual direction of translation. However, half the test set is made up of translations on the source side. The translationese effect means that the model will usually get higher scores on this half of the test set, potentially inflating the score. Consequently, the intrinsic metrics could suggest choosing a model that does not actually perform well on the desired task (translating *from* native text).

We investigate this effect in our own work by examining the mean BLEU scores for each model on each half of the test sets, giving the results in Figure 5. Each bar indicates the mean percentage change in BLEU scores over the parallel-only baseline model for the models trained on the different BT datasets, so a larger bar means a better performing model. The left-hand bars in each quadrant show the performance of each model on the back-translated half of the test set (*to native*) and the right-hand bars give the performance of each model on the forward-translated half of the test set (*from native*).

We see a significant translationese effect for all models, as the percentage change in scores over the baseline are much higher when the models translate already translated text (the left-hand side bars are higher than the right-hand ones). However, it appears that the *nucleus* dataset is less affected by the translationese effect than the other datasets, since it shows less of a decline in performance when translating native text. This may be due to a similar regularising effect as discussed previously, as it is more difficult for the model to overfit to BT data when it is generated with nucleus sampling. A direction for future research is how to obtain the benefits of using monolingual data (as BT does) without exacerbating the translationese effect.

5 Related work

Improving back translation The original paper introducing BT by Sennrich et al. (2016) found that using a higher-quality NMT system for BT led to higher BLEU scores in the final trained system. This finding was corroborated by Burlot and Yvon (2018), and following work has investigated further ways to improve NMT. These include iterative BT (Hoang et al., 2018), targeting difficult words (Fadaee and Monz, 2018), and tagged BT (Caswell et al., 2019). Section 3.2.1 of Haddow et al. (2021)




Figure 5: The mean percentage change in BLEU score for each model on the test set(s) over the parallel-only models, separated by language direction. The left-hand side (*to native*) has translated text on the source side and native text on the target side of the test set (back translation). The right-hand side (*from native*) has native text on the source side and translated text on the target side of the test set.

presents a comprehensive survey of BT and its variants as applied to low-resource NMT.

Diversity in machine translation Most of the work on the lack of diversity in machine-translated text are in the context of automatic evaluation (Edunov et al., 2020; Graham et al., 2020; Roberts et al., 2020). As for diversity in BT specifically, Edunov et al. (2018) argue that MAP prediction, as is typically used to generate BT through beam search, leads to overly-regular synthetic source sentences which do not cover the true data distribution. They propose instead generating BT with sampling or noised beam outputs, and find model performance increases for all but the lowest resource scenarios. Alternatively, Soto et al. (2020) generate diverse BT by training multiple machine-translation systems with varying architectures.

Generating diversity Increasing diversity in BT is part of the broader field of diverse generation, by which we mean methods to vary the surface form of a production whilst keeping the meaning as similar as possible. Aside from generating diverse translations (Gimpel et al., 2013; He et al., 2018; Shen et al., 2019; Nguyen et al., 2020; Li et al., 2021), it is also used in question answering systems (Sultan et al., 2020), visually-grounded generation (Vi-

jayakumar et al., 2018), conversation models (Li et al., 2016), and particularly paraphrasing (Mallinson et al., 2017; Wieting and Gimpel, 2018; Hu et al., 2019; Thompson and Post, 2020; Goyal and Durrett, 2020; Krishna et al., 2020). Some recent work such as Iyyer et al. (2018), Huang and Chang (2021), and Hosking and Lapata (2021) explicitly model the meaning and the form of the input separately. In this way, they aim to vary the syntax of the output whilst preserving the semantics so as to generate more diverse paraphrases. Unfortunately, these methods are difficult to apply to a low-resource scenario as they require external resources (e.g. accurate syntactic parsers, large-scale paraphrase data) which are not available for most of the world’s languages.

6 Conclusion

In this paper, we introduced a two-part framework for understanding diversity in NMT data, splitting it into *lexical diversity* and *syntactic diversity*. Our empirical analysis suggests that whilst high amounts of both types of diversity are important in training data, lexical diversity may be more beneficial than syntactic. In addition, achieving high diversity in BT should not be at the expense of ad-

equacy. We find that generating BT with nucleus sampling results in the highest final NMT model performance for our systems. Future work could investigate further the affect of high lexical diversity on BT independent of syntactic diversity.

Acknowledgements

This work was supported in part by the UKRI Centre for Doctoral Training in Natural Language Processing, funded by the UKRI (grant EP/S022481/1) and the University of Edinburgh, School of Informatics and School of Philosophy, Psychology & Language Sciences. It was also supported by funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 825299 (GoURMET) and funding from the UK Engineering and Physical Sciences Research Council (EPSRC) fellowship grant EP/S001271/1 (MTStretch).

The experiments in this paper were performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/P020259/1), and DiRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk).

Finally, the authors would like to thank our anonymous reviewers for their time and helpful comments, and we give special thanks to Henry Conklin and Bailin Wang for their help with tree kernels and many useful discussions.

References

Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa, Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, Jungo Kasai, Daniel Khashabi, Kevin Knight, Tom Kocmi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydryn, and Marcos Zampieri. 2021. [Findings of the 2021 conference on machine translation \(WMT21\)](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88, Online. Association for Computational Linguistics.

Marta Bañón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Esplà-Gomis, Mikel L. Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, Sergio Ortiz Rojas, Leopoldo Pla Sempere, Gema Ramírez-Sánchez, Elsa Sarrías, Marek Strelec, Brian Thompson, William Waites, Dion Wiggins, and Jaume Zaragoza. 2020. [ParaCrawl: Web-scale acquisition of parallel corpora](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4555–4567, Online. Association for Computational Linguistics.

Starkaður Barkarson and Steinþór Steingrímsson. 2019. [Compiling and filtering ParLce: An English-Icelandic parallel corpus](#). In *Proceedings of the 22nd Nordic Conference on Computational Linguistics*, pages 140–145, Turku, Finland. Linköping University Electronic Press.

Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R. Costa-jussa, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(WMT20\)](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1–55, Online. Association for Computational Linguistics.

Lukas Biewald. 2020. [Experiment tracking with weights and biases](#). Software available from wandb.com.

Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Philipp Koehn, and Christof Monz. 2018. [Findings of the 2018 conference on machine translation \(WMT18\)](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 272–303, Belgium, Brussels. Association for Computational Linguistics.

Franck Burlot and François Yvon. 2018. [Using monolingual data in neural machine translation: a systematic study](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 144–155, Brussels, Belgium. Association for Computational Linguistics.

Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. [Tagged back-translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63, Florence, Italy. Association for Computational Linguistics.

Danqi Chen and Christopher Manning. 2014. [A fast and accurate dependency parser using neural networks](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 740–750, Doha, Qatar. Association for Computational Linguistics.

Henry Conklin, Bailin Wang, Kenny Smith, and Ivan Titov. 2021. [Meta-learning to compositionally gen-](#)

- eralize. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3322–3335, Online. Association for Computational Linguistics.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Marzieh Fadaee and Christof Monz. 2018. [Back-translation sampling by targeting difficult words in neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 436–446, Brussels, Belgium. Association for Computational Linguistics.
- Kevin Gimpel, Dhruv Batra, Chris Dyer, and Gregory Shakhnarovich. 2013. [A systematic exploration of diversity in machine translation](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1100–1111, Seattle, Washington, USA. Association for Computational Linguistics.
- Tanya Goyal and Greg Durrett. 2020. [Neural syntactic reordering for controlled paraphrase generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 238–252, Online. Association for Computational Linguistics.
- Yvette Graham, Christian Federmann, Maria Eskevich, and Barry Haddow. 2020. [Assessing human-parity in machine translation on the segment level](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4199–4207, Online. Association for Computational Linguistics.
- Barry Haddow, Rachel Bawden, Antonio Valerio Miceli Barone, Jindrich Helcl, and Alexandra Birch. 2021. [Survey of low-resource machine translation](#). *CoRR*, abs/2109.00486.
- Xuanli He, Gholamreza Haffari, and Mohammad Norouzi. 2018. [Sequence to sequence mixture model for diverse machine translation](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 583–592, Brussels, Belgium. Association for Computational Linguistics.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. [Iterative back-translation for neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 18–24, Melbourne, Australia. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text de-generation](#). In *International Conference on Learning Representations*.
- Tom Hosking and Mirella Lapata. 2021. [Factorising meaning and form for intent-preserving paraphrasing](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1405–1418, Online. Association for Computational Linguistics.
- J Edward Hu, Rachel Rudinger, Matt Post, and Benjamin Van Durme. 2019. [Parabank: Monolingual bitext generation and sentential paraphrasing via lexically-constrained neural machine translation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 6521–6528.
- Kuan-Hao Huang and Kai-Wei Chang. 2021. [Generating syntactically controlled paraphrases without using annotated parallel pairs](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1022–1033, Online. Association for Computational Linguistics.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. [Adversarial example generation with syntactically controlled paraphrase networks](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1875–1885, New Orleans, Louisiana. Association for Computational Linguistics.
- Haukur Páll Jónsson, Haukur Barri Símonarson, Vésteinn Snæbjarnarson, Steinþór Steingrímsson, and Hrafn Loftsson. 2020. [Experimenting with different machine translation models in medium-resource settings](#). In *International Conference on Text, Speech, and Dialogue*, pages 95–103. Springer.
- Kalpesh Krishna, John Wieting, and Mohit Iyyer. 2020. [Reformulating unsupervised style transfer as paraphrase generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 737–762, Online. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [SentencePiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.

- Jicheng Li, Pengzhi Gao, Xuanfu Wu, Yang Feng, Zhongjun He, Hua Wu, and Haifeng Wang. 2021. [Mixup decoding for diverse machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 312–320, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Jonathan Mallinson, Rico Sennrich, and Mirella Lapata. 2017. [Paraphrasing revisited with neural machine translation](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 881–893, Valencia, Spain. Association for Computational Linguistics.
- Alessandro Moschitti. 2006. [Making tree kernels practical for natural language learning](#). In *11th Conference of the European Chapter of the Association for Computational Linguistics*, pages 113–120, Trento, Italy. Association for Computational Linguistics.
- Xuan-Phi Nguyen, Shafiq Joty, Kui Wu, and Ai Ti Aw. 2020. Data diversification: A simple strategy for neural machine translation. *Advances in Neural Information Processing Systems*, 33:10018–10029.
- Myle Ott, Michael Auli, David Grangier, and Marc’Aurelio Ranzato. 2018. [Analyzing uncertainty in neural machine translation](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3956–3965. PMLR.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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 of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Ricardo Rei, 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. Association for Computational Linguistics.
- Nicholas Roberts, Davis Liang, Graham Neubig, and Zachary C. Lipton. 2020. [Decoding and diversity in machine translation](#). In *Advances in Neural Information Processing Systems*, volume 33. Curran Associates, Inc.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. [WikiMatrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Tianxiao Shen, Myle Ott, Michael Auli, and Marc’Aurelio Ranzato. 2019. [Mixture models for diverse machine translation: Tricks of the trade](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5719–5728. PMLR.
- Raphael Shu, Hideki Nakayama, and Kyunghyun Cho. 2019. [Generating diverse translations with sentence codes](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1823–1827, Florence, Italy. Association for Computational Linguistics.
- Xabier Soto, Dimitar Shterionov, Alberto Poncelas, and Andy Way. 2020. [Selecting backtranslated data from multiple sources for improved neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3898–3908, Online. Association for Computational Linguistics.
- Steinþór Steingrímsson, Sigrún Helgadóttir, Eiríkur Rögnvaldsson, Starkaður Barkarson, and Jón Guðnason. 2018. [Risamálheild: A very large Icelandic text corpus](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

- Md Arafat Sultan, Shubham Chandel, Ramón Fernandez Astudillo, and Vittorio Castelli. 2020. [On the importance of diversity in question generation for QA](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5651–5656, Online. Association for Computational Linguistics.
- Guy Tevet and Jonathan Berant. 2021. [Evaluating the evaluation of diversity in natural language generation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 326–346, Online. Association for Computational Linguistics.
- Brian Thompson and Matt Post. 2020. [Paraphrase generation as zero-shot multilingual translation: Disentangling semantic similarity from lexical and syntactic diversity](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 561–570, Online. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Eva Vanmassenhove, Dimitar Shterionov, and Andy Way. 2019. [Lost in translation: Loss and decay of linguistic richness in machine translation](#). In *Proceedings of Machine Translation Summit XVII: Research Track*, pages 222–232, Dublin, Ireland. European Association for Machine Translation.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ashwin Vijayakumar, Michael Cogswell, Ramprasaath Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. 2018. [Diverse beam search for improved description of complex scenes](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- John Wieting and Kevin Gimpel. 2018. [ParaNMT-50M: Pushing the limits of paraphrastic sentence embeddings with millions of machine translations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 451–462, Melbourne, Australia. Association for Computational Linguistics.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Texygen: A benchmarking platform for text generation models](#). In *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '18*, page 1097–1100, New York, NY, USA. Association for Computing Machinery.

A Appendix

A.1 Representative examples from back-translated datasets (translated from Icelandic)

Original *Pjóðverjar hafa tekið forræðið og stefnt er að stofnun stórríkis.*

Beam

- The Germans have taken custody and are aimed at the establishment of a large state.
- The Germans have taken custody and are aimed at the creation of a large state.
- The Germans have taken custody and are aimed at establishing a large state.

Comment: Only one or two words differ between sentences (underlined).

Sampling

- The Germz governmentregluru has committed suicide, intending to organise a major state.
- The Germano had ensured that British commanders in France would be aides of theærd rapidly.
- And the need to defend and establish theseUC-tions are all organized intomissions from Ir-aqéttihe.

Comment: Sentences show large variation in structure and vocabulary, but they contain many non-dictionary words (underlined) and adequacy is low.

Nucleus

- Germany has taken custody and aimed to establish a large country.
- The German government initiated a group operation, to establish capital city.
- The Germany has managed to make an example of their full widowed demands.

Comment: There is a moderate amount of variation between sentences in terms of syntax and vocabulary, but no non-dictionary words. Some phrases lack adequacy (underlined).

Syntax-groups

- The Germans have taken custody and are aimed at the establishment of a large state.
- The Icelandic Institute of Natural History
- As a result, the Germans have taken control of the country and are aimed at establishing a large state.

Comment: The second and third sentences contain hallucinations, presumably in order to generate according to the syntactic templates (underlined).

A.2 Model architecture and infrastructure

All NMT models in this paper are transformer models (Vaswani et al., 2017). We conducted a hyper-parameter search for each language pair, training English↔Turkish and English↔Icelandic NMT models and using the BLEU score as the optimisation metric. We give the settings which differ to *transformer-base* in Table 5. We use the same hyper-parameter settings for all models trained for the same language pair.

We use the Fairseq toolkit to train all our NMT models (Ott et al., 2019). We train on four NVIDIA A100-SXM-80GB GPUs and use CUDA 11.1 plus a Python 3.8 Conda environment provided in the Github repo. We generate on one GPU, since to our knowledge the Fairseq toolkit does not support multi-GPU decoding. We use Weights and Biases for experiment tracking (Biewald, 2020).

	tr-en	is-en
Dropout	0.6	0.3
Activation dropout	0.1	0
Attention dropout		0.1
Learning rate		0.001
L.R. scheduler	Inv. square root	
Optimiser	Adam	
Optimiser parameters	0.9, 0.98	
Label smoothing	0.1	
Shared embeddings	all	
Batch size	64	
Update frequency	16	
Patience	15	

Table 5: Hyper-parameter settings for NMT transformer models trained for each language pair. All other settings are the default for *transformer-base*.