

# The Effectiveness of Morphology-aware Segmentation in Low-Resource Neural Machine Translation

**Jonne Sälevä**

Brandeis University

jonesaleva@brandeis.edu

**Constantine Lignos**

Brandeis University

lignos@brandeis.edu

## Abstract

This paper evaluates the performance of several modern subword segmentation methods in a low-resource neural machine translation setting. We compare segmentations produced by applying BPE at the token or sentence level with morphologically-based segmentations from LMVR and MORSEL. We evaluate translation tasks between English and each of Nepali, Sinhala, and Kazakh, and predict that using morphologically-based segmentation methods would lead to better performance in this setting. However, comparing to BPE, we find that no consistent and reliable differences emerge between the segmentation methods. While morphologically-based methods outperform BPE in a few cases, what performs best tends to vary across tasks, and the performance of segmentation methods is often statistically indistinguishable.

## 1 Introduction

Despite the advances of neural machine translation (NMT), building effective translation systems for lower-resourced and morphologically rich languages remains a challenging process. The lack of large training data sets tends to lead to problems of vocabulary sparsity, a problem exacerbated by the combinatorial explosion of permissible surface forms commonly encountered when working with morphologically rich languages.

Current NMT systems typically operate at the level of *subwords*. Most commonly, these systems achieve vocabulary reduction by decomposing tokens into character sequences constructed by maximizing an information-theoretic compression criterion. The most widely used subword segmentation method is byte pair encoding, originally invented in the data compression literature by Gage (1994), and introduced to the MT community by Sennrich et al. (2016). Another approach to open vocabulary

NMT has been to compose characters or character n-grams to form word representations (Ataman and Federico, 2018a; Ling et al., 2015).

As BPE has become mainstream, the question of whether segmenting words in a linguistically-informed fashion provides a benefit remains open. Intuitively, the translation task may be easier when using subwords that contain maximal linguistic signal, as opposed to heuristically derived units based on data compression. The greatest benefit may come in low-resource settings, where the training data is small and biases toward morphological structure may lead to more reusable units.

We seek to address this question by exploring the usefulness of linguistically-motivated subword segmentation methods in NMT, as measured against a BPE baseline. Specifically, we investigate the effectiveness of morphology-based segmentation algorithms of Ataman et al. (2017) and Lignos (2010) as alternatives to BPE at the word or sentence level and find that they do not lead to reliable improvements under our experimental conditions. We perform our evaluation using both BLEU (Papineni et al., 2002) and CHRF3 (Popović, 2015). In our low-resource NMT setting, all these methods provide comparable results.

The contribution of this work is that it provides insights into the performance of these segmentation methods using a thorough experimental paradigm in a highly replicable environment. We evaluate without the many possible confounds related to back-translation and other processes used in state-of-the-art NMT systems, focusing on the performance of a straightforward Transformer-based system. To analyze the performance differences between the various segmentation strategies, we utilize a Bayesian linear model as well as nonparametric hypothesis tests.

Translation task	Split	Sentences	Tokens (EN)	Tokens (non-EN)
NE ↔ EN	Train	563,947	4,483,440	4,200,818
SI ↔ EN	Train	646,781	4,837,496	4,180,520
KK ↔ EN	Train (120k)	124,770	379,546	319,484
KK ↔ EN	Train (220k)	222,424	1,717,414	1,365,605
NE ↔ EN	Dev	2,559	46,267	37,576
SI ↔ EN	Dev	2,898	53,471	48,659
KK ↔ EN	Dev	2,066	45,975	37,258
NE ↔ EN	Test	2,835	51,455	43,802
SI ↔ EN	Test	2,766	50,973	46,318
KK → EN	Test	1,000	20,376	15,943
EN → KK	Test	998	24,074	19,141

Table 1: Number of sentences in raw corpora. The 120k and 220k training conditions for KK correspond to training KK↔EN models with/without an additional crawled corpus. The test sets for KK→EN and EN→KK are different from each other and mirror the released WMT19 data.

## 2 Related work

Attempts to create unsupervised, morphologically-aware segmentations have often been derived from the Morfessor family of morphological segmentation tools (Virpioja et al., 2013). In addition to extensions of Morfessor, such as *Cognate Morfessor* (Grönroos et al., 2018), Ataman et al. (2017) and Ataman and Federico (2018b) introduced the LMVR model, derived from Morfessor *FlatCat* (Grönroos et al., 2014), and applied it to NMT tasks on Arabic, Czech, German, Italian, Turkish and English, noting that LMVR outperforms a BPE baseline in CHRF3 and BLEU. Contrary to their results, however, Toral et al. (2019) find that using LMVR yielded mixed results: on a Kazakh-English translation task the authors observed marginal BLEU improvements over BPE, whereas for English-Kazakh, the authors reported LMVR to perform marginally worse than BPE in terms of CHRF3.

There have also been efforts to combine BPE with linguistically motivated approaches. For instance, Huck et al. (2017) propose to combine BPE with various linguistic heuristics such as prefix, suffix, and compound splitting. The authors work with English-German and German-English tasks, and observe performance improvements of approximately 0.5 BLEU compared to a BPE-only baseline. As another example, Weller-Di Marco and Fraser (2020) combine BPE with a full morphological analysis on the source and target sides of an English-German translation task, and report performance improvements exceeding 1 BLEU point

over a BPE-only baseline.

Finally, even though Sennrich et al. (2016) originally only used the NMT training set to train their segmentation model, others have recently found benefit in adding monolingual data to the process. In particular, Scherrer et al. (2020) used both SentencePiece and Morfessor as segmentation models on an Upper Sorbian-German translation task and found a monotonic increase in BLEU when the segmentation model was trained with additional data, while at the same time keeping the NMT training data constant.

## 3 Experiments

To investigate the effect of subword segmentation algorithms on NMT performance, we train translation models using the Transformer architecture of Vaswani et al. (2017). We base our work on two recent datasets: FLoRes (Guzmán et al., 2019), and select languages from the WMT 2019 Shared Task on News Translation (Barrault et al., 2019). Corpus statistics for all corpora can be found in Table 1.

The FLoRes dataset consists of two language pairs, English-Nepali and English-Sinhala. To add another lower-resourced language, we use the Kazakh-English translation data from WMT19. In terms of morphological typology, both Nepali and Sinhala are agglutinative languages (Prasain, 2011; Priyanga et al., 2017), as is Kazakh (Kessikbayeva and Cicekli, 2014).

We conduct two sets of experiments on Kazakh to investigate how the amount of training data influences our results: first, we train only on

Segmentation	Sentence
Original	The nation slowly started being centralized and during
SentencePiece	_the _n ation _sl ow ly _start ed _being _cent ral ized _and _d ur ing
Subword-NMT	the n@@ ation s@@ low@@ ly star@@ ted being cen@@ tr@@ ali@@ z@@ ed and d@@ ur@@ ing
LMVR	the nation s +low +ly st +ar +ted be +ing c +ent +ral +ized and d +ur +ing
MORSEL	the nation s@@ low +ly start +ed being cen@@ tr@@ ali@@ z +ed and du@@ r +ing

Table 2: Examples of segmentation strategies and tokenization.

the WikiTitles and News Commentary corpora (train120k), followed by another set of experiments (train220k) where we include the web crawl corpus prepared by Bagdat Myrzakhmetov of Nazarbayev University. We also conducted experiments with Gujarati data from WMT19, but BLEU scores were too low to allow for meaningful analysis. For our models, we generally follow the architecture and hyperparameter choices of the FLoRes Transformer baseline, except for setting `clip_norm` to 0.1 and enabling FP16 training.

Despite the widespread use of auxiliary techniques such as back-translation we deliberately refrain from employing such techniques in this work. This is done to best isolate the effect of varying the subword segmentation algorithm, and to avoid the complexity of disentangling it from the effect of other factors. It should be noted, however, that such techniques were highly prevalent among of systems submitted to the KK $\leftrightarrow$ EN WMT19 News Translation Shared Task: 64% used back-translation, 61% used ensembling, and 57% employed extensive corpus filtering (Barrault et al., 2019).

### 3.1 Subword segmentation algorithms

Below we describe our hyperparameter settings for the various subword segmentation algorithms. Sinhala and Nepali are tokenized using the Indic NLP tokenizer (Kunchukuttan, 2020), whereas for English and Kazakh we use the Moses tokenizer (Koehn et al., 2007). Example segmentations from actual data can be seen in Table 2.

The segmentation methods we evaluate learn their subword vocabularies from frequency distributions of tokenized text. The exception to this is SentencePiece, whose subword units are learned from sentences, including whitespace. In the case of English and Kazakh, these sentences are untokenized whereas for Nepali and Sinhala, preprocessing with the Indic NLP tokenizer is applied following the approach of Guzmán et al. (2019).

#### 3.1.1 Subword-NMT and SentencePiece

As our baseline subword segmentation algorithm, we use the BPE implementation from Subword-NMT<sup>1</sup>. Throughout our experiments we use a joint vocabulary of the source and target and set the number of requested symbols to 5,000. For SentencePiece, we use the default BPE implementation<sup>2</sup> with a joint vocabulary size of 5,000 words. These choices are motivated by the general observation by Sennrich and Zhang (2019) that lowering BPE size improves translation quality in ultra-low resource conditions, and the specific value of 5,000 was previously used by Guzmán et al. (2019). The same small vocabulary size has been used elsewhere in the low-resource NMT literature, for instance by Roest et al. (2020) while training NMT systems for Inuktitut. We also conducted a hyperparameter sweep for 2,500, 5,000, 7,500 and 10,000 merge operations, but noticed no improvement over the choice of 5,000 motivated by prior work.

#### 3.1.2 LMVR

For LMVR (Ataman et al., 2017), we utilize slightly modified versions of the sample scripts from the author’s Github repository<sup>3</sup>. Our main modification is tuning the `corpusweight` hyperparameter in the Morfessor Baseline (Virpioja et al., 2013) model used to seed the LMVR model. Tuning is performed by maximizing the F1 score for segmenting the English side of the training data, using the English word lists from the Morpho Challenge 2010 shared task (Kurimo et al., 2010) as gold standard segmentations. After tuning the Morfessor Baseline model, we train a separate LMVR model for each language in a language pair using a vocabulary size parameter of 2,500 per language.

<sup>1</sup><https://github.com/rsennrich/subword-nmt>

<sup>2</sup><https://github.com/google/sentencepiece>

<sup>3</sup><https://github.com/d-ataman/lmvr>

### 3.1.3 MORSEL

MORSEL (Lignos, 2010) provides linguistically-motivated unsupervised morphological analysis that has been shown to work effectively on small datasets (Chan and Lignos, 2010). While it provides derivations of morphologically complex forms via a combination of stems and affix rules, we modified it to provide a segmentation and then postprocessed its output to apply BPE to the stems to yield a limited-size vocabulary.

For example, on the English side of the NE-EN training data, MORSEL analyzes the word *algebraic* as resulting from the stem *algebra* being combined with the suffix rule *+ic*. A BPE model is trained on all of the stems in MORSEL’s analysis, and when that is applied to the stem, it is segmented as `al@@ ge@@ br@@ a`. The stem and suffix are combined using a special plus character to denote suffixation, so the final segmentation is `al@@ ge@@ br@@ a +ic`. Tuning is performed as with LMVR, using the English word lists from the Morpho Challenge 2010 shared task (Kurimo et al., 2010) as a reference. We adjust the number of BPE units learned from the stems to keep the total per-language vocabulary below 2,500.

## 4 Results and analysis

Our experimental results can be seen in Table 3. All BLEU scores were computed using `sacrebleu`, and all CHRF3 scores using `nltk`. Each row consists of the mean and standard deviation computed across 5 random seeds for each configuration. We also plot the raw results in Figure 1. Table 4 gives counts for the number of times each segmentation approach was the top-performing one or statistically indistinguishable from it. Table 7 in the appendix gives  $p$ -values for all comparisons performed.

Overall, based on Tables 3 and 4, no segmentation method seems to emerge as the clear winner across translation tasks, although BPE applied at the token (Subword-NMT) or sentence (SentencePiece) level performs well consistently. Subword-NMT or SentencePiece perform best in 12 out of 16 cases (counting BLEU and CHRF3 for each translation task), while morphology-based methods rank best in 4 out of 16 cases. In particular, we note that morphology-based methods seem to achieve or tie the best BLEU performance for translation tasks involving SI, and best CHRF3 performance for

Segm. method	BLEU	CHRF3
EN-KK (train120k)		
LMVR	1.00 ± 0.12	<b>21.98</b> ± 0.41
MORSEL	0.94 ± 0.11	<b>21.24</b> ± 0.89
SentencePiece	1.04 ± 0.09	<b>21.48</b> ± 0.47
Subword-NMT	<u>1.32</u> ± 0.08	<u>22.12</u> ± 0.28
EN-KK (train220k)		
LMVR	1.82 ± 0.13	<b>22.74</b> ± 0.84
MORSEL	<b>2.06</b> ± 0.11	<b>22.88</b> ± 0.40
SentencePiece	<u>2.18</u> ± 0.08	<b>22.78</b> ± 0.43
Subword-NMT	1.94 ± 0.22	<b>22.62</b> ± 0.88
KK-EN (train120k)		
LMVR	1.70 ± 0.07	23.72 ± 0.44
MORSEL	<b>2.62</b> ± 0.08	<u>26.26</u> ± 0.36
SentencePiece	2.34 ± 0.21	24.64 ± 0.81
Subword-NMT	<u>3.14</u> ± 0.18	<b>25.92</b> ± 0.54
KK-EN (train220k)		
LMVR	9.42 ± 0.26	33.88 ± 0.76
MORSEL	<b>10.44</b> ± 0.48	<b>34.58</b> ± 0.88
SentencePiece	<b>10.02</b> ± 0.29	33.50 ± 0.54
Subword-NMT	<u>10.68</u> ± 0.34	<u>35.52</u> ± 0.41
EN-NE		
LMVR	4.32 ± 0.04	<b>31.00</b> ± 0.29
MORSEL	<b>4.38</b> ± 0.16	<b>31.28</b> ± 0.47
SentencePiece	<u>4.58</u> ± 0.15	<u>31.36</u> ± 0.35
Subword-NMT	<b>4.42</b> ± 0.16	30.96 ± 0.34
NE-EN		
LMVR	7.84 ± 0.11	<b>34.10</b> ± 0.16
MORSEL	5.30 ± 0.30	28.18 ± 0.97
SentencePiece	<b>8.42</b> ± 0.23	<b>34.40</b> ± 0.73
Subword-NMT	<u>8.46</u> ± 0.15	<b>34.18</b> ± 0.13
EN-SI		
LMVR	<u>1.44</u> ± 0.32	<b>28.22</b> ± 0.30
MORSEL	<b>1.12</b> ± 0.13	27.44 ± 0.34
SentencePiece	<b>1.08</b> ± 0.31	<b>27.56</b> ± 0.43
Subword-NMT	0.88 ± 0.13	26.78 ± 0.51
SI-EN		
LMVR	7.24 ± 0.22	32.16 ± 0.63
MORSEL	<b>7.78</b> ± 0.16	<b>34.32</b> ± 0.30
SentencePiece	<b>7.52</b> ± 0.08	<b>33.58</b> ± 0.43
Subword-NMT	<b>7.76</b> ± 0.25	<u>34.38</u> ± 0.38

Table 3: Mean and standard deviation of BLEU and CHRF3 across translation tasks and segmentation methods. Underlined values represent the highest mean scores. Bolded values are not significantly different ( $p > 0.05$ ) than the highest score as determined by Dunn’s test.

KK-EN with smaller training data (`train120k`) as well as EN-SI. However, when using LMVR, we fail to find the significant gains in BLEU compared to BPE reported by Ataman et al. (2017).

Comparing our results to Guzmán et al. (2019), we note that the scores are similar, although not directly comparable as we report lowercased BLEU




Figure 1: CHRF3 vs. BLEU with different translation tasks indicated by color and segmentation by marker shape.

Segmentation method	BLEU	CHRF3
Subword-NMT	6	6
SentencePiece	6	6
MORSEL	6	6
LMVR	1	5

Table 4: Number of times each segmentation method was or tied with the best-performing method under each metric, counted across all tasks.

scores.<sup>4</sup> They report EN-NE/NE-EN baseline BLEU scores of 4.3 and 7.6 using a single random seed, which are in line with our results in Table 3. For EN-SI/SI-EN, the authors report 1.2 and 7.2 BLEU, which likewise matches our findings. Even though our scores are low overall, they are as low as is to be expected using this approach, size of data, and languages. In order to compare our results to WMT19 participant systems, it is only meaningful to compare our system to baseline systems due to the widespread use of auxiliary training techniques, such as back-translation. For instance, Casas et al. (2019) report baseline NMT scores of 2.32 on KK-EN and 1.42 on EN-KK, which are in line with our MORSEL and SentencePiece results on KK-EN, and Subword-NMT results on EN-KK in the train120k condition.

#### 4.1 Modeling BLEU and CHRF3

Based on Figure 1 and Tables 3 and 4, the BLEU and CHRF3 scores vary with both the translation task and segmentation method. Intuitively, the

<sup>4</sup>We lowercased all data in preprocessing because MORSEL and Morfessor, which LMVR is derived from, are designed to operate on lowercase inputs.

Pairwise comparison	$\tau$ (BLEU)	$\tau$ (CHRF3)
SentencePiece - Subword-NMT	$-0.05 \pm 0.08$	$-0.07 \pm 0.20$
MORSEL - Subword-NMT	$-0.12 \pm 0.07$	$0.02 \pm 0.18$
LMVR - Subword-NMT	$-0.26 \pm 0.06$	$-0.19 \pm 0.21$

Table 5: Posterior means and standard deviations of  $\tau_m - \tau_{\text{Subword-NMT}}$  (pairwise comparison with BPE) under the BLEU and CHRF3 models. Values are rounded to two decimal places.

scores seem to cluster around a certain range for each translation task, and are perturbed slightly depending on the choice of segmentation method. To better disentangle the influence of these factors, we fit a Bayesian linear model to the experimental data, treating the final BLEU/CHRF3 score as a sum of a “translation task effect”  $\eta$ , a “segmentation method effect”  $\tau$ , and a translation task-specific noise term  $\epsilon$ .<sup>5</sup> The  $\eta$  and  $\epsilon$  terms are estimated for each of the eight translation tasks (e.g. SI-EN and EN-SI are estimated separately), and  $\tau$  is estimated for each of the four segmentation methods using results from all translation tasks.

To explicitly compare SentencePiece, LMVR and MORSEL to the Subword-NMT baseline, we also model the pairwise differences between each method’s  $\tau$ -term and that of Subword-NMT. The posterior inferences for these quantities can be seen in Table 5 and are plotted in the appendix. For BLEU, the differences for LMVR are several standard deviations below 0, suggesting that it performs worse than the Subword-NMT baseline

<sup>5</sup>In the appendix, Section A gives details of our model, and Table 6 gives the point estimates of the posterior mean and standard deviation for  $\eta$  and  $\tau$ .

when accounting for all translation tasks. Similarly, MORSEL is almost 2 standard deviations away from 0, though its posterior interval does cover 0. In both cases, the effect size is small, with a mean of -0.12 and -0.26 points of BLEU for MORSEL and LMVR, respectively. The reliability of this difference also disappears for LMVR under the CHR3 model, where no segmentation method’s posterior mean is several standard deviations away from 0.

We hypothesize that this greater discrimination among methods when using BLEU may originate from the differences between how BLEU and CHR3 operate. Since CHR3 is a character-level metric, it is less prone than BLEU to penalizing a given translation due to subword outputs that are *almost* correct. For instance, consider output of `do@@ gs` → `dogs` with `dog` as the reference; while CHR3 awards credit for this as a partial match, BLEU treats it as entirely incorrect. This further underscores our observation that segmentation methods perform inconsistently across experimental conditions.

## 5 Conclusion and future work

Contrary to our hypothesis about the usefulness of morphology-aware segmentation, we see no consistent advantage, and possibly a small disadvantage, to using LMVR or MORSEL in this resource-constrained setting. By and large, our experiments and modeling show that no segmentation approach consistently achieves the best BLEU/CHR3 across all translation tasks. BPE remains a good default segmentation strategy, but it is possible that LMVR, MORSEL, or similar systems may show larger performance advantages for languages with specific morphological structures.

Consequently, we believe further work is needed to better understand when morphology-aware methods are most effective and to develop methods that provide a consistent advantage over BPE. One such avenue of future work would be to broaden our analysis to more languages and include languages that are higher-resourced but morphologically rich and as well as ones that are lower-resourced but morphologically poor. [Ortega et al. \(2021\)](#), which we encountered during preparation of the final version of this paper, began to address these questions by comparing Morfessor with BPE and their own BPE variant on Finnish, Quechua and Spanish.

An alternative approach which we intend to pur-

sue in future work is experimenting with supervised morphological segmenters or analyzers that can be efficiently developed even in lower-resourced settings. Incorporating such “gold standard” segmentations may make it clearer whether the unsupervised morphological segmenters are capturing linguistically-relevant structure.

Finally, there is the question of whether BPE can approximate a general representation for a language instead of converging on a corpus-specific set of subwords. To test this, one can add monolingual data and train the BPE segmentation on that larger data set. Ideally the new, “enriched” segmentations would depend less on the specific vocabulary of the training corpus. As noted above, [Scherrer et al. \(2020\)](#) observed this approach to be helpful in terms of BLEU. However, it remains unknown why the subwords derived from a larger corpus perform better, and whether better identification of morphological structure could be responsible.

We hope that this work and these ideas will catalyze further research, and that efficient methods for translating to and from lower-resourced languages can be developed as a result.

## References

- Duygu Ataman and Marcello Federico. 2018a. [Compositional representation of morphologically-rich input for neural machine translation](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 305–311, Melbourne, Australia. Association for Computational Linguistics.
- Duygu Ataman and Marcello Federico. 2018b. [An evaluation of two vocabulary reduction methods for neural machine translation](#). In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 97–110, Boston, MA. Association for Machine Translation in the Americas.
- Duygu Ataman, Matteo Negri, Marco Turchi, and Marcello Federico. 2017. [Linguistically motivated vocabulary reduction for neural machine translation from turkish to english](#). *The Prague Bulletin of Mathematical Linguistics*, 108(1):331–342.
- Loïc Barrault, 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. Association for Computational Linguistics.
- Noe Casas, José A. R. Fonollosa, Carlos Escolano, Christine Basta, and Marta R. Costa-jussà. 2019. The TALP-UPC machine translation systems for WMT19 news translation task: Pivoting techniques for low resource MT. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 155–162, Florence, Italy. Association for Computational Linguistics.
- Erwin Chan and Constantine Lignos. 2010. Investigating the relationship between linguistic representation and computation through an unsupervised model of human morphology learning. *Research on Language and Computation*, 8(2-3):209–238.
- Philip Gage. 1994. A new algorithm for data compression. *C Users Journal*, 12(2):23–38.
- Stig-Arne Grönroos, Sami Virpioja, and Mikko Kurimo. 2018. Cognate-aware morphological segmentation for multilingual neural translation. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 386–393, Belgium, Brussels. Association for Computational Linguistics.
- Stig-Arne Grönroos, Sami Virpioja, Peter Smit, and Mikko Kurimo. 2014. Morfessor FlatCat: An HMM-based method for unsupervised and semi-supervised learning of morphology. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1177–1185, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6098–6111, Hong Kong, China. Association for Computational Linguistics.
- Matthias Huck, Simon Riess, and Alexander Fraser. 2017. Target-side word segmentation strategies for neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pages 56–67, Copenhagen, Denmark. Association for Computational Linguistics.
- Gulshat Kessikbayeva and Ilyas Cicekli. 2014. Rule based morphological analyzer of Kazakh language. In *Proceedings of the 2014 Joint Meeting of SIGMORPHON and SIGFSM*, pages 46–54, Baltimore, Maryland. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Anoop Kunchukuttan. 2020. The IndicNLP Library. [https://github.com/anoopkunchukuttan/indic\\_nlp\\_library/blob/master/docs/indicnlp.pdf](https://github.com/anoopkunchukuttan/indic_nlp_library/blob/master/docs/indicnlp.pdf).
- Mikko Kurimo, Sami Virpioja, Ville T Turunen, et al. 2010. Proceedings of the Morpho Challenge 2010 workshop. In *Morpho Challenge Workshop 2010*. Aalto University School of Science and Technology.
- Constantine Lignos. 2010. Learning from unseen data. In *Proceedings of the Morpho Challenge 2010 Workshop*, pages 35–38.
- Wang Ling, Isabel Trancoso, Chris Dyer, and Alan W Black. 2015. Character-based neural machine translation. *ArXiv*.
- John E. Ortega, Richard Castro Mamani, and Kyunghyun Cho. 2021. Neural machine translation with a polysynthetic low resource language. *Machine Translation*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing 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.
- Balaram Prasain. 2011. *Computational analysis of Nepali morphology: A model for natural language processing*. Ph.D. thesis, Faculty of Humanities and Social Sciences of Tribhuvan University.
- Rajith Priyanga, Surangika Ranatunga, and Gihan Dias. 2017. Sinhala word joiner. In *Proceedings of the 14th International Conference on Natural Language Processing (ICON-2017)*, pages 220–226, Kolkata, India. NLP Association of India.
- Christian Roest, JK Spenader, and A Toral Ruiz. 2020. Morphological segmentation of polysynthetic languages for neural machine translation: The case of inuktitut. mastersthesis, University of Groningen.

Yves Scherrer, Stig-Arne Grönroos, and Sami Virpioja. 2020. [The University of Helsinki and aalto university submissions to the WMT 2020 news and low-resource translation tasks](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 1129–1138, Online. Association for Computational Linguistics.

Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

Rico Sennrich and Biao Zhang. 2019. [Revisiting low-resource neural machine translation: A case study](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 211–221, Florence, Italy. Association for Computational Linguistics.

Antonio Toral, Lukas Edman, Galiya Yeshmagambetova, and Jennifer Spenser. 2019. [Neural machine translation for English–Kazakh with morphological segmentation and synthetic data](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 386–392, Florence, Italy. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.

Sami Virpioja, Peter Smit, Stig-Arne Grönroos, Mikko Kurimo, et al. 2013. Morfessor 2.0: Python implementation and extensions for Morfessor baseline. *Technical Report*.

Marion Weller-Di Marco and Alexander Fraser. 2020. [Modeling word formation in English–German neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4227–4232, Online. Association for Computational Linguistics.

## A Bayesian Linear Model Details

Mathematically, our model can be expressed as:

$$\phi_{lm} = \eta_l + \tau_m + \epsilon_l \quad (1)$$

where  $\phi_{lm} \in \{\text{BLEU}, \text{CHRF3}\}$ ,  $\eta_l$  and  $\tau_m$  represent the “translation task effect” and “segmentation method effect,” and  $\epsilon_l$  is a translation task-specific variance term.

To initialize our Bayesian linear model from Equation 1, we set the following priors. For the BLEU model,  $\eta_l \sim \mathcal{N}(4, 3)$  and  $\tau_m \sim \mathcal{N}(0, 1)$ .

Segmentation method effect	$\tau$ (BLEU)	$\tau$ (CHRF3)
LMVR	$-0.09 \pm 0.47$	$0.41 \pm 0.50$
MORSEL	$0.05 \pm 0.47$	$0.63 \pm 0.50$
SentencePiece	$0.12 \pm 0.47$	$0.53 \pm 0.50$
Subword-NMT	$0.17 \pm 0.47$	$0.60 \pm 0.50$
Pairwise comparison	$\tau$ (BLEU)	$\tau$ (CHRF3)
SentencePiece - Subword-NMT	$-0.05 \pm 0.08$	$-0.07 \pm 0.20$
LMVR - Subword-NMT	$-0.26 \pm 0.06$	$-0.19 \pm 0.21$
MORSEL - Subword-NMT	$-0.12 \pm 0.07$	$0.02 \pm 0.18$
Translation task effect	$\eta$ (BLEU)	$\eta$ (CHRF3)
EN-KK (train120k)	$1.01 \pm 0.47$	$21.16 \pm 0.52$
EN-KK (train220k)	$1.94 \pm 0.47$	$22.21 \pm 0.51$
EN-NE	$4.36 \pm 0.47$	$30.60 \pm 0.50$
EN-SI	$1.07 \pm 0.47$	$26.95 \pm 0.52$
KK-EN (train120k)	$2.39 \pm 0.48$	$24.58 \pm 0.56$
KK-EN (train220k)	$10.07 \pm 0.48$	$33.81 \pm 0.54$
NE-EN	$7.41 \pm 0.56$	$32.02 \pm 0.82$
SI-EN	$7.51 \pm 0.47$	$33.05 \pm 0.54$

Table 6: Posterior means and standard deviations for  $\tau$  and  $\eta$  under the BLEU and CHRF3 models.

For the CHRF3 model,  $\eta_l \sim \mathcal{N}(15, 7)$  and  $\tau_m \sim \mathcal{N}(0, 1)$ . The priors are the same regardless of translation task or segmentation method. For our noise terms, we use a  $\epsilon_l \sim \text{HalfCauchy}(5)$  prior in all models. Our rationale for these priors is that  $\eta_l$  should place most of its probability mass within the observed range of BLEU/CHRF3, whereas  $\tau_m$  should, *a priori*, take on positive and negative values with equal probability, reflecting a lack of prior information. All models are fit using PyMC3, and MCMC posterior inference performed using the No-U-Turn Sampler.

All posterior means for  $\eta$  are close to the average BLEU/CHRF3 scores per translation task observed in Table 3, and fall between 1.01 and 10.07 for the BLEU model, and 21.16 and 33.81 for the CHRF3 model. In contrast, the posterior means for  $\tau$  are universally small:  $-0.09$ ,  $0.05$ ,  $0.12$ , and  $0.17$  for LMVR, MORSEL, SentencePiece and Subword-NMT, respectively, with a posterior standard deviation of  $0.47$ . The  $\tau$ -terms under the CHRF3 model exhibit a similar pattern:  $0.41$ ,  $0.63$ ,  $0.53$ ,  $0.60$ , with a posterior standard deviation of  $0.50$ . Compared to the posterior standard deviation, as well as translation task effects  $\eta$ , the  $\tau$ -terms are practically 0. This, in conjunction with our analysis using Dunn’s test, suggests that there is not a segmentation method that consistently works best across translation tasks.

Figures 2 and 3 show posterior predictive distributions for the BLEU and CHRF3 models. Figure 4 shows the posterior distribution of pairwise differences between each of the other segmentation methods and Subword-NMT.



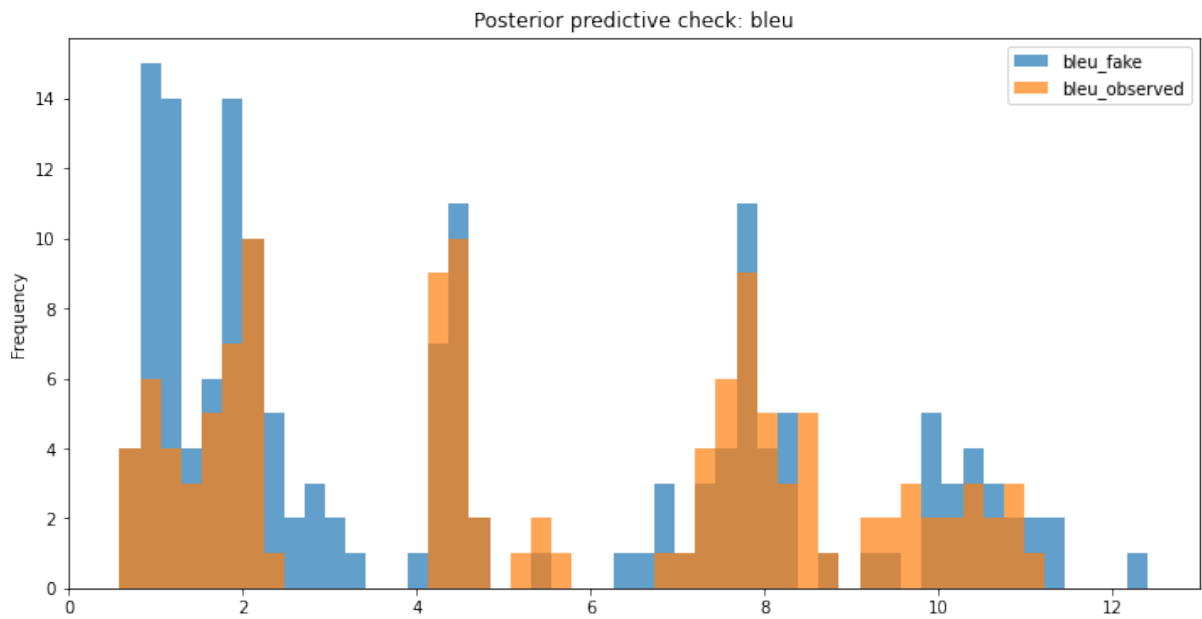


Figure 2: Posterior predictive distribution of BLEU under the Bayesian linear model.




Figure 3: Posterior predictive distribution of CHRF3 under the Bayesian linear model.

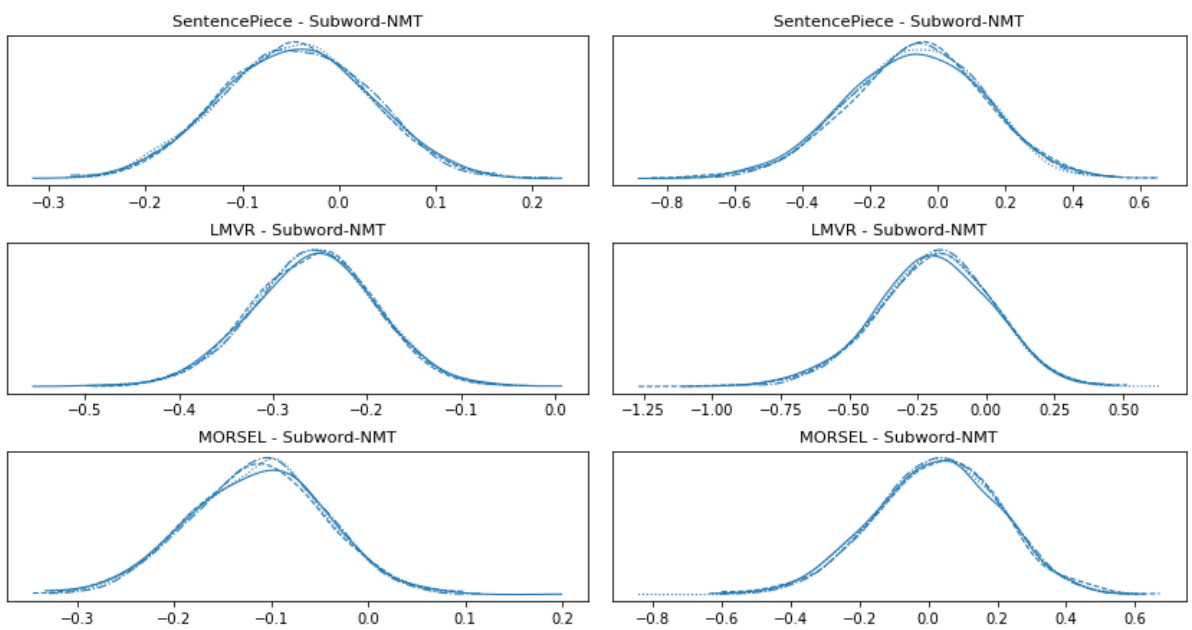


Figure 4: Posterior distribution of pairwise differences  $\tau_m - \tau_{\text{Subword-NMT}}$  in the BLEU model (left) and CHRF3 model (right). Note:  $m \in \{\text{SentencePiece, LMVR, MORSEL}\}$

Language pair	Segmentation method	$p$ -value (BLEU)	$p$ -value (CHRF3)
EN-NE	LMVR	<b>0.014</b>	0.071
EN-NE	MORSEL	0.057	0.466
EN-NE	SentencePiece	1.000	1.000
EN-NE	Subword-NMT	0.137	<b>0.046</b>
NE-EN	LMVR	<b>0.036</b>	0.405
NE-EN	MORSEL	<b>0.001</b>	<b>0.002</b>
NE-EN	SentencePiece	0.872	1.000
NE-EN	Subword-NMT	1.000	0.767
EN-SI	LMVR	1.000	1.000
EN-SI	MORSEL	0.246	<b>0.036</b>
EN-SI	SentencePiece	0.071	0.091
EN-SI	Subword-NMT	<b>0.003</b>	<b>0.000</b>
SI-EN	LMVR	<b>0.002</b>	<b>0.001</b>
SI-EN	MORSEL	1.000	0.851
SI-EN	SentencePiece	0.080	0.057
SI-EN	Subword-NMT	0.850	1.000
KK-EN (train220k)	LMVR	<b>0.001</b>	<b>0.009</b>
KK-EN (train220k)	MORSEL	0.592	0.149
KK-EN (train220k)	SentencePiece	0.069	<b>0.002</b>
KK-EN (train220k)	Subword-NMT	1.000	1.000
EN-KK (train220k)	LMVR	<b>0.002</b>	0.788
EN-KK (train220k)	MORSEL	0.216	0.768
EN-KK (train220k)	SentencePiece	1.000	1.000
EN-KK (train220k)	Subword-NMT	<b>0.037</b>	0.893
KK-EN (train120k)	LMVR	<b>0.000</b>	<b>0.001</b>
KK-EN (train120k)	MORSEL	0.140	1.000
KK-EN (train120k)	SentencePiece	<b>0.011</b>	<b>0.026</b>
KK-EN (train120k)	Subword-NMT	1.000	0.611
EN-KK (train120k)	LMVR	<b>0.008</b>	0.872
EN-KK (train120k)	MORSEL	<b>0.001</b>	0.068
EN-KK (train120k)	SentencePiece	<b>0.032</b>	0.096
EN-KK (train120k)	Subword-NMT	1.000	1.000

Table 7: Dunn’s test  $p$ -values for BLEU and CHRF3. Boldface indicates statistical significance at the  $\alpha = 0.05$  level.