

A Systematic Analysis of Subwords and Cross-Lingual Transfer in Multilingual Translation

Francois Meyer and Jan Buys

Department of Computer Science

University of Cape Town

francois.meyer@uct.ac.za, jbuys@cs.uct.ac.za

Abstract

Multilingual modelling can improve machine translation for low-resource languages, partly through shared subword representations. This paper studies the role of subword segmentation in cross-lingual transfer. We systematically compare the efficacy of several subword methods in promoting synergy and preventing interference across different linguistic typologies. Our findings show that subword regularisation boosts synergy in multilingual modelling, whereas BPE more effectively facilitates transfer during cross-lingual fine-tuning. Notably, our results suggest that differences in orthographic word boundary conventions (the morphological granularity of written words) may impede cross-lingual transfer more significantly than linguistic unrelatedness. Our study confirms that decisions around subword modelling can be key to optimising the benefits of multilingual modelling.

1 Introduction

Machine translation (MT) models have become increasingly multilingual (Dabre et al., 2020). This greatly benefits low-resource languages through positive transfer from high-resource languages (Ha et al., 2016; Aharoni et al., 2019). However, increasing multilinguality in a limited shared parameter space can lead to suboptimal performance for high-resource languages (Firat et al., 2016; Team et al., 2022). There is a tradeoff between maximising positive cross-lingual transfer (also known as synergy) while minimising negative cross-lingual interaction (also known as interference).

Several modelling decisions affect synergy and interference in multilingual MT. Shaham et al. (2023) experimentally analysed the influence of factors like model size and language data proportions. One aspect their study failed to consider is subword segmentation. The shared subword vocabulary of multilingual models presents a similar

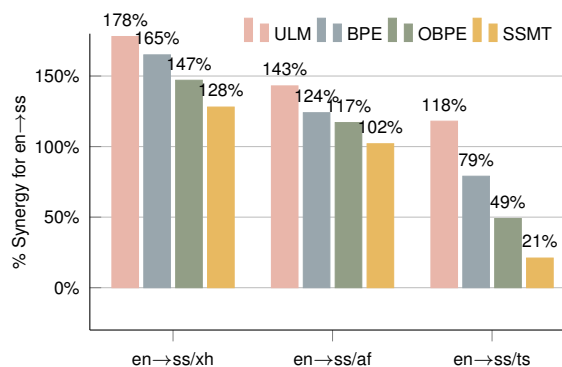


Figure 1: Performance increase for English→Siswati through multilingual modelling varies greatly across subword methods and linguistic contexts.

trade-off as the shared parameter space - overlapping subword representations induce synergy, but having to represent multiple languages in a limited vocabulary can harm cross-lingual transfer (Chung et al., 2020; Rust et al., 2021; Patil et al., 2022).

In this paper we experimentally analyse the role of subword segmentation in multilingual and cross-lingual MT. Our goal is to compare different classes of subword methods with regards to their ability to induce synergy, reduce interference, and transfer knowledge during cross-lingual finetuning. We also investigate how cross-lingual transfer is influenced by the linguistic similarities of interacting languages, with particular focus on factors related to subword structure like morphological typology and orthographic word boundary conventions (the degree to which morphemes are concatenated or written as separate orthographic words).

We run experiments on translation from English to four linguistically diverse South African languages (see Table 1). This selection covers different levels of language relatedness, morphological complexity, and orthographic word granularity, allowing us to analyse how these factors interact with different subword methods to influence cross-

Language	Family	Morphology	Orthography	<i>What is your name?</i>	<i>Thank you!</i>
Siswati (ss)	NC/Bantu/Nguni	agglutinative	conjunctive	Ngubani ligama lakho?	Ngiyabonga!
isiXhosa (xh)	NC/Bantu/Nguni	agglutinative	conjunctive	Ungubani igama lakho?	Enkosi!
Setswana (ts)	NC/Bantu/Sotho-Tswana	agglutinative	disjunctive	Leina la gago ke mang?	Ke a leboga!
Afrikaans (af)	Indo-European/Germanic	analytic	disjunctive	Wat is jou naam?	Dankie!

Table 1: We vary the language modelled alongside Siswati to control relatedness, morphology, and orthography.

lingual transfer. Low-resource languages stand to benefit most from multilingual modelling. In all our experiments we focus on cross-lingual transfer to Siswati, which is by far the least resourced among the languages included. It presents exactly the type of real world low-resource translation scenario we are interested in studying.

We conduct two sets of experiments - multilingual MT and cross-lingual finetuning. Our multilingual experiments follow [Shaham et al. \(2023\)](#) in training several trilingual MT models and comparing synergy/interference (see Figure 1). In the cross-lingual finetuning experiments we finetune pretrained bilingual MT models on new languages. Our results demonstrate that decisions around subword segmentation significantly affect MT performance. ULM ([Kudo, 2018](#)) improves synergy in multilingual modelling, while BPE ([Sennrich et al., 2016](#)) enhances cross-lingual transfer during finetuning. Going beyond linguistic relatedness, we find that the much less studied influence of orthographic word boundary conventions can drastically affect the cross-lingual transfer achieved between interacting languages.

2 Related Work

Synergy and interference are well-established phenomena ([Firat et al., 2016](#); [Aharoni et al., 2019](#); [Team et al., 2022](#)), but not well understood. [Shaham et al. \(2023\)](#) address this by systematically analysing the role of several factors in synergy and interference: (1) model size, (2) data size, (3) language proportions, (4) number of languages, and (5) language relatedness. Their results show that scaling model size and tuning the data sampling temperature greatly alleviates interference. They do not vary subword segmentation in their experiments, using the same sentencepiece ([Kudo and Richardson, 2018](#)) vocabulary across all models.

However, multilingual vocabularies are known to affect cross-lingual transfer through factors such as cross-lingual subword overlap ([Pires et al., 2019](#); [Wu and Dredze, 2019](#); [Patil et al., 2022](#)) and under-

represented low-resource languages ([Wang et al., 2021](#); [Ács, 2019](#)). These issues have mainly been studied for multilingual language modelling ([Rust et al., 2021](#); [Maronikolakis et al., 2021](#); [Chung et al., 2020](#)), but the same concerns hold for MT ([Wang et al., 2020a](#)). We are unaware of existing work comparing the multitude of proposed subword methods in the context of multilingual MT.

3 Methodology

This study involves two sets of MT experiments - (1) multilingual (trilingual) experiments to investigate synergy/interference, and (2) finetuning experiments to analyse cross-lingual transfer. Our goal is to determine which subwords benefit low-resource languages and how cross-lingual transfer depends on linguistic typology. The linguistic diversity of South Africa is an ideal testing ground for our purposes. Siswati is a low-resource agglutinative language, so effective subword modelling is critical for dealing with the inevitably high proportion of out-of-vocabulary words in held-out datasets. We use Siswati as the low-resource target language in our experiments and alternate the higher resourced language modelled alongside Siswati between isiXhosa, Setswana, and Afrikaans.

Table 1 shows how these language present varying linguistic relationships to Siswati. IsiXhosa is closely related. Setswana is somewhat related and also agglutinative, but diverges in its orthography - its writing system is disjunctive ([Pretorius et al., 2009](#)). This refers to how linguistic words (e.g. nouns, verbs) are represented as orthographic words (space-separated tokens). Disjunctive orthographies write a single linguistic word as multiple orthographic words (e.g. in Setswana prefixal morphemes are space-separated from verbal roots).

While linguistic relatedness and morphological complexity are obvious features to consider in any analysis of cross-lingual interactions, we are unaware of work considering the impact of orthographic word boundary conventions. We include it as a factor in our study because of its potential

relevance to subword segmentation. Orthographic word boundaries determine the pre-tokenisation of text before subword segmenters are applied, so it could well affect aspects like segmentation granularity and overlap between the subword vocabularies of different languages. Afrikaans is linguistically unrelated to Siswati and also disjunctive, but because of its analytic morphology (lower morpheme-to-word ratio) its written words are sometimes more aligned to those of Siswati (e.g. see phrases in Table 1).

This selection of languages allows us to isolate the cross-lingual effects of linguistic relatedness, morphological typology, and orthographic word boundary conventions. In the case of Setswana-Siswati, we can study whether the potentially positive cross-lingual effect of their linguistic relatedness is nullified by the fact that the two languages have very different conventions for orthographic word boundaries.

3.1 Multilingual Modelling

We train two bilingual models and one trilingual model per language pair (see Table 2).

Languages	Examples	Subwords
en→ss	166k	BPE/ULM/SSMT
en→xh/ts/af	1.6m	BPE/ULM/SSMT
en→ss+xh/ts/af	1.6m+166k	BPE/ULM/SSMT/OBPE

Table 2: Multilingual experimental setup: bilingual and trilingual models (bilingual OBPE is equivalent to BPE).

This setup allows us to compare how MT performance changes for en→ss and en→xh/ts/af going from bilingual models to multilingual models. Following Shaham et al. (2023), we measure synergy/interference for a translation direction $s \rightarrow t$ by the relative difference in performance between a bilingual model trained to translate only from s to t and a multilingual model trained to translate to an additional language.

Shaham et al. (2023) use test set cross-entropy loss to measure MT performance, but this cannot be reliably used to compare across different subword segmentations. Instead, we use test set chrF++ (Popović, 2017) to measure performance. It is a reference-based metric that combines word and character information, so it is well suited for evaluating subword-level performance. Our modified

formula for measuring synergy/interference is

$$\mathcal{I}_{s \rightarrow t} = \frac{\text{CHR}F_{++}(M_{s \rightarrow t}^{\text{multi}}) - \text{CHR}F_{++}(M_{s \rightarrow t}^{\text{bi}})}{\text{CHR}F_{++}(M_{s \rightarrow t}^{\text{bi}})},$$

where M are trained multilingual/bilingual models evaluated on $s \rightarrow t$ translation. Negative values of $\mathcal{I}_{s \rightarrow t}$ indicate worse performance for $s \rightarrow t$ in the multilingual model (interference) and positive values indicate improved performance (synergy).

3.2 Cross-Lingual Finetuning

We train a bilingual subword segmenter and MT model for en→xh/ts/af, and then finetune and evaluate the model in the other translation directions (e.g. pretrain en→xh and finetune on en→ss, en→ts, and en→af). Varying the subword method reveals how different subwords facilitate cross-lingual transfer during finetuning from higher resourced languages (isiXhosa/Setswana/Afrikaans) to lower resourced Siswati.

4 Experimental Setup

We compare five subword segmenters (four per experiment). We chose methods that represent the main paradigms of subword segmentation - deterministic segmentation, subword regularisation, learning subwords during training, and subword techniques for enhancing cross-lingual transfer.

- BPE** (Sennrich et al., 2016) iteratively adds frequently co-occurring subwords to its vocabulary. We use it as a deterministic segmenter.
- ULM** (Kudo, 2018) learns segmentation to optimise a unigram language model and can be used as a probabilistic segmenter, exposing models to multiple subword segmentations for regularisation. We set the sampling parameter α to 0.5.
- SSMT** (Meyer and Buys, 2023) is a subword segmental MT model which learns subword segmentation jointly during MT training, with the goal of learning subwords that optimise MT performance.
- OBPE** (Patil et al., 2022) modifies BPE to boost subword overlap among languages in multilingual vocabularies. We use it in our multilingual experiments to see if increased shared subword representations promote synergy.
- XBPE** (Wang et al., 2020b) extends the BPE vocabulary of a pretrained model to include BPE subwords of a new translation direction. New subword embeddings are randomly initialised. We use it in our finetuning experiments to see if the vocabulary extension enhances cross-lingual transfer.

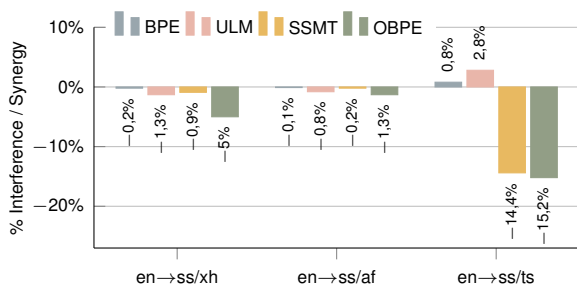


Figure 2: Performance change for en→xh/af/ts through multilingual modelling alongside en→ss.

Model	tgt	BPE	ULM	SSMT	OBPE
en→ss/xh	ss	33.4	35.1	33.6	31.1
	xh	46.8	47.2	46.1	44.6
en→ss/af	ss	28.2	30.6	29.7	27.4
	af	60.9	61.1	60.2	60.2
en→ss/ts	ss	22.5	27.5	17.8	18.8
	ts	31.3	34.7	23.6	26.4

Table 3: Test set chrF++ of trilingual models.

Training We train models on WMT22 data (Ade-lani et al., 2022) and validate and test on FLORES (Goyal et al., 2021; Team et al., 2022). The number of training sentences are shown in Table 2. For en→ss this is the full WMT22 dataset, but for en→xh/ts/af we sampled sentences from en→xh and en→ts to match the size of en→af, removing data size as an influence. We also removed examples from en→xh/ts/af where English source sentences were found in en→ss to neutralise the positive transfer effect of multi-parallel overlap (Stap et al., 2023). The hyperparameters of our models and subword methods are included in Appendix A.

5 Results & Discussion

We plot the synergy/interference analysis of our multilingual experiments in Figures 1 & 2, while the absolute performance of the models are provided in Table 3. The results from our cross-lingual finetuning experiments are presented in Figure 3.

5.1 Which subwords promote synergy and minimise interference?

ULM consistently achieves greater synergy than other subword methods. This holds across all linguistic contexts (Fig. 1) and results in better absolute performance in all translation directions (Table 3). It comes at the cost of minimal interference for the higher resourced languages, and even some synergy for en→ts (Fig. 2). The subword regulari-

		Finetuned and evaluated on en →							
		xh	af	ts	ss	xh	af	ts	ss
Pretrained on en ↑	xh	46.9	60.4	38.7	32.3	47.8	15.8	1.9	31.1
	af	46.4	61.0	36.9	28.3	13.5	61.6	17.3	28.2
	ts	43.8	58.5	31.1	25.5	10.8	3.6	33.8	24.7
		(a) BPE				(b) ULM			
		xh	af	ts	ss	xh	af	ts	ss
Pretrained on en ↑	xh	46.5	59.9	38.0	32.2	46.9	59.5	37.3	29.9
	af	46.6	60.3	36.3	30.1	46.1	61.0	35.6	25.0
	ts	42.0	56.8	27.6	23.4	43.4	58.3	31.1	21.9
		(c) SSMT				(d) XBPE			

Figure 3: Test set chrF++ of pretraining for en→xh/af/ts (rows) and finetuning on en→xh/af/ts/ss (columns). Diagonal entries are bilingual models with no finetuning.

sation of ULM ensures that models are more robust to the varied subwords of multilingual modelling.

5.2 Which subwords transfer cross-lingually?

BPE subwords exhibit the greatest cross-lingual transferability. In contrast to our multilingual findings, the subword regularisation of ULM proves a barrier to cross-lingual finetuning. ULM is a probabilistic segmenter that is sampled during training, but when the probabilistic model is based on one language and applied to another, its samples might yield highly inadequate subword units. The consistent deterministic segmentation of BPE allows the finetuned model to adapt to a new translation direction effectively.

5.3 What is the role of linguistic typology?

A consistent pattern emerges in the cross-lingual dynamics between Siswati and other languages. IsiXhosa modelling proves to be most beneficial for Siswati performance. Afrikaans achieves less transfer, presumably because it is not related. Somewhat surprisingly, the weakest synergy is between Siswati and Setswana, even though both are agglutinative Bantu languages. This highlights the impact of orthographic systems on cross-lingual transfer: *diverging word boundary conventions can impede cross-lingual transfer more than linguistic unrelatedness*. Data-driven multilingual models that learn from text might miss underlying similarities between languages that are obscured by superficial differences in their surface realisations.

Our results highlight two interacting effects. Firstly, linguistic relatedness does play a role –

isiXhosa consistently improves Siswati more than Setswana and Afrikaans. Secondly, in the specific case of Setswana-Siswati, their relatedness is rendered all but irrelevant by the fact that the languages have diverging orthographies. Afrikaans does not have the extremely disjunctive orthography of Setswana so even though it is less related to Siswati than Setswana, the orthography of Setswana prevents transfer to Siswati. Linguistic distance plays a role in both cases (Afrikaans-Siswati and Setswana-Siswati) but for Setswana-Siswati it is a less important factor than orthography.

Orthography is a notable difference between Nguni languages like Siswati and Sotho-Tswana languages like Setswana (Pretorius et al., 2009). Taljard and Bosch (2006) showed that the diverging orthographies of these two language groups necessitate different approaches to more traditional NLP tasks, even though the languages are linguistically and morphologically related. Our results suggest a similar situation for cross-lingual transfer in multilingual modelling: Differences in orthographic word boundary conventions harms synergy between otherwise related languages.

6 Conclusion

We presented an in-depth study on the role of subwords in multilingual and cross-lingual MT. Our results demonstrate that subword segmentation significantly influences cross-lingual interactions. ULM proves optimal for transfer to low-resource languages in multilingual modelling, while BPE enables greater cross-lingual transfer during finetuning. Besides language relatedness, we show that similarities/differences in orthographic word granularity can greatly affect multilingual performance. There is more work to be done on the role of orthographic word boundary conventions in neural MT. Future work could aim to design multilingual techniques that see past orthographic differences in order to leverage more fundamental similarities between languages.

Limitations

Our study is limited to translation from English to four South African languages. While the chosen languages are typologically diverse, our conclusions might not necessarily hold for languages from different language families and with distinct orthographies. We did not consider languages that

have multiple orthographies, which might be another approach to study the effects of orthography. The performance differences between different subword segmentation methods across languages in our results are relatively consistent, but a more detailed analysis on the interaction between choice of subword segmentation method and language could yield additional explanations of the results.

Acknowledgements

This work is based on research supported in part by the National Research Foundation of South Africa (Grant Number: 129850). Computations were performed using facilities provided by the University of Cape Town’s ICTS High Performance Computing team: hpc.uct.ac.za.

References

- David Adelani, Md Mahfuz Ibn Alam, Antonios Anastopoulos, Akshita Bhagia, Marta R. Costa-jussà, Jesse Dodge, Fahim Faisal, Christian Federmann, Natalia Fedorova, Francisco Guzmán, Sergey Koshelev, Jean Maillard, Vukosi Marivate, Jonathan Mbuya, Alexandre Mourachko, Safiyyah Saleem, Holger Schwenk, and Guillaume Wenzek. 2022. [Findings of the WMT’22 shared task on large-scale machine translation evaluation for African languages](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 773–800, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Roe Aharoni, Melvin Johnson, and Orhan Firat. 2019. [Massively multilingual neural machine translation](#). 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 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hyung Won Chung, Dan Garrette, Kiat Chuan Tan, and Jason Riesa. 2020. [Improving multilingual models with language-clustered vocabularies](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4536–4546, Online. Association for Computational Linguistics.
- Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. [A survey of multilingual neural machine translation](#). *ACM Comput. Surv.*, 53(5).
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016. [Multi-way, multilingual neural machine translation with a shared attention mechanism](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 866–875, San

- Diego, California. Association for Computational Linguistics.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2021. The flores-101 evaluation benchmark for low-resource and multilingual machine translation.
- Thanh-Le Ha, Jan Niehues, and Alex Waibel. 2016. [Toward multilingual neural machine translation with universal encoder and decoder](#). In *Proceedings of the 13th International Conference on Spoken Language Translation*, Seattle, Washington D.C. International Workshop on Spoken Language Translation.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. 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.
- Antonis Maronikolakis, Philipp Dufter, and Hinrich Schütze. 2021. [Wine is not v i n. on the compatibility of tokenizations across languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2382–2399, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Francois Meyer and Jan Buys. 2023. [Subword segmental machine translation: Unifying segmentation and target sentence generation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2795–2809, Toronto, Canada. Association for Computational Linguistics.
- Vaidehi Patil, Partha Talukdar, and Sunita Sarawagi. 2022. [Overlap-based vocabulary generation improves cross-lingual transfer among related languages](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 219–233, Dublin, Ireland. Association for Computational Linguistics.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- Rigardt Pretorius, Ansu Berg, Laurette Pretorius, and Biffie Viljoen. 2009. [Setswana tokenisation and computational verb morphology: Facing the challenge of a disjunctive orthography](#). In *Proceedings of the First Workshop on Language Technologies for African Languages*, pages 66–73, Athens, Greece. Association for Computational Linguistics.
- Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. [How good is your tokenizer? on the monolingual performance of multilingual language models](#). 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 3118–3135, 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.
- Uri Shaham, Maha Elbayad, Vedanuj Goswami, Omer Levy, and Shruti Bhosale. 2023. [Causes and cures for interference in multilingual translation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15849–15863, Toronto, Canada. Association for Computational Linguistics.
- David Stap, Vlad Niculae, and Christof Monz. 2023. [Viewing knowledge transfer in multilingual machine translation through a representational lens](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14973–14987, Singapore. Association for Computational Linguistics.
- Elsabé Taljard and Sonja E. Bosch. 2006. [A comparison of approaches to word class tagging: Disjunctively vs. conjunctively written bantu languages](#). *Nordic Journal of African Studies*, 15:428–442.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semafor Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#).
- Changhan Wang, Kyunghyun Cho, and Jiatao Gu. 2020a. [Neural machine translation with byte-level subwords](#). In *AAAI 2020 - 34th AAAI Conference on*

Artificial Intelligence, AAAI 2020 - 34th AAAI Conference on Artificial Intelligence, pages 9154–9160. AAAI press. Publisher Copyright: Copyright © 2020 Association for the Advancement of Artificial Intelligence. All rights reserved.; 34th AAAI Conference on Artificial Intelligence, AAAI 2020 ; Conference date: 07-02-2020 Through 12-02-2020.

Xinyi Wang, Sebastian Ruder, and Graham Neubig. 2021. [Multi-view subword regularization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 473–482, Online. Association for Computational Linguistics.

Zihan Wang, Karthikeyan K, Stephen Mayhew, and Dan Roth. 2020b. [Extending multilingual BERT to low-resource languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2649–2656, Online. Association for Computational Linguistics.

Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). 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 833–844, Hong Kong, China. Association for Computational Linguistics.

Judit Ács. 2019. [Exploring BERT’s vocabulary](#).

A Model Configurations

We use the model size and training hyperparameters of the fairseq [transformer-base](#) architecture. We train our models for 45 epochs and use a language sampling temperature of $T = 1.5$ to balance exposure to low-resource and high-resource languages. In our cross-lingual finetuning experiments we finetune models on en→ss for 20 epochs and on en→xh/af/ts for 10 epochs.

When applying subword methods we specify a shared vocabulary size of 8k. This is slightly smaller than the optimal vocabulary size of 10k used by [Meyer and Buys \(2023\)](#) in experiments on these same languages, since we use smaller subsets of the WMT22 datasets.