



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2021

How Suitable Are Subword Segmentation Strategies for Translating Non-Concatenative Morphology?

Amrhein, Chantal ; Sennrich, Rico

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-208885>

Conference or Workshop Item

Published Version



The following work is licensed under a Creative Commons: Attribution-NonCommercial 4.0 International (CC BY-NC 4.0) License.

Originally published at:

Amrhein, Chantal; Sennrich, Rico (2021). How Suitable Are Subword Segmentation Strategies for Translating Non-Concatenative Morphology? In: Findings of the Association for Computational Linguistics: EMNLP 2021, Punta Cana, Dominican Republic, 7 November 2021 - 11 November 2021. ACL Anthology, 689-705.

How Suitable Are Subword Segmentation Strategies for Translating Non-Concatenative Morphology?

Chantal Amrhein¹ and Rico Sennrich^{1,2}

¹Department of Computational Linguistics, University of Zurich

²School of Informatics, University of Edinburgh

{amrhein, sennrich}@cl.uzh.ch

Abstract

Data-driven subword segmentation has become the default strategy for open-vocabulary machine translation and other NLP tasks, but may not be sufficiently generic for optimal learning of non-concatenative morphology. We design a test suite to evaluate segmentation strategies on different types of morphological phenomena in a controlled, semi-synthetic setting. In our experiments, we compare how well machine translation models trained on subword- and character-level can translate these morphological phenomena. We find that learning to analyse and generate morphologically complex surface representations is still challenging, especially for non-concatenative morphological phenomena like reduplication or vowel harmony and for rare word stems. Based on our results, we recommend that novel text representation strategies be tested on a range of typologically diverse languages to minimise the risk of adopting a strategy that inadvertently disadvantages certain languages.¹

1 Introduction

Data-driven subword-level segmentation of text (Sennrich et al., 2016; Kudo, 2018) is a well-known and widely used text representation strategy in the natural language processing (NLP) community. While subword segmentation largely solves the open vocabulary problem, previous research has shown that models often break down in out-of-domain contexts (El Boukkouri et al., 2020), when encountering spelling errors (Belinkov and Bisk, 2018; Pruthi et al., 2019), when translating morphologically-rich languages (Ataman and Federico, 2018) and in multilingual scenarios (Chung et al., 2020; Wang et al., 2021). The reason for this is that even slight deviations from the text seen

¹Test suite and code available at <https://github.com/ZurichNLP/segtest>

when learning a segmentation model can result in entirely different segmentations and often aggressively over-segmented text.

Given the rich morphological diversity across natural languages, it is especially interesting to investigate the suitability of subword segmentation to represent different morphological phenomena. For example, reduplication is a non-concatenative morphological phenomenon² that is common across the world’s languages, but is marginal in higher-resource European languages,³ which raises the question if the dominant text representation strategies inadvertently disadvantage NLP systems for languages that feature it.

Many types of morphological phenomena (see examples in Table 1) pose challenges to subword-level models. For concatenative phenomena such as affixes, subword-level segmentations often do not adhere to morpheme boundaries which can hurt the performance of these models. For non-concatenative morphology, it is still unclear to what extent subword-level models can learn to generalise to rare or unseen words. We believe these challenges are exciting opportunities to work on better text representations for cross-lingual NLP but currently, there is a lack of targeted evaluation environments. Most previous work evaluates very specific morphological or morpho-syntactic functions such as number, case, gender or subject-verb agreement (Sennrich, 2017; Burlot and Yvon, 2017; Warstadt et al., 2020) rather than evaluating how well different *types* of morphological phenomena can be learned.

To address this issue, we design a test suite that can be used to evaluate how well a range of morphological phenomena can be learned with sentence-level sequence-to-sequence models. We focus on

²The distinction between non-concatenative and concatenative morphological phenomena is often a topic of debate in Linguistics. We follow Lieber and Štekauer (2014).

³See reduplication feature in WALS (Rubino, 2013).

Compounding (German)	Circumfixation (Chickasaw)	Infixation (Bontoc)	Vowel Harmony (Turkish)	Reduplication (Itza’)
Schild / Kröte ‘shield’ / ‘toad’	lakna ‘it is yellow’	fikas ‘strong’	üzüldünüz üz-ül-dü-nüz	tz’eek ‘few’
Schildkröte ‘turtle’	iklakno ‘it isn’t yellow’ (Fromkin et al., 2018)	fumikas ‘to be strong’ (Fromkin et al., 2018)	sadden-PASS-PAST-2PL ‘You became sad.’ (Göksel and Kerslake, 2005)	tz’eek-tz’eek ‘very few’ (Hoffing, 2000)

Table 1: Examples for the morphological phenomena studied in this paper. The first three are concatenative, the last two non-concatenative.

the task of neural machine translation (NMT) and evaluate in a semi-synthetic DE→EN setting, allowing for an automatic evaluation that controls for various confounding factors. In our experiments, we test how well current segmentation strategies on subword- and character-level can learn to translate compounds, circumfixed words, infixed words, vowel harmony and reduplicated words.

Our contributions are the following:

- We design an evaluation environment for various types of morphological phenomena that can be used to evaluate future text representation strategies.
- We find that non-concatenative morphological phenomena and generalisation to rare word bases are especially challenging to learn with current segmentation strategies.
- We show that subword segmentation is less suitable to learn the correct surface form but all segmentation strategies perform well when we represent the morphological phenomena with an abstract token instead.

2 Related Work

Isolated morphological analysis and inflection have long been of interest to the NLP community, with yearly shared tasks (Kurimo et al., 2010; Vylomova et al., 2020) that result in dedicated architectures that perform well for many languages (Aharoni and Goldberg, 2017; Makarov and Clematide, 2018; Wu and Cotterell, 2019; Rios et al., 2021). However, despite the large morphological diversity of natural languages, many approaches for sentence-level sequence-to-sequence tasks are often only tested on a subset of (morphologically similar) languages and then adopted without much questioning (Bender, 2011).

One such example is subword-level representation of text (Sennrich et al., 2016; Kudo, 2018) which has contributed greatly to the success of deep learning in various NLP tasks and has become a necessary preprocessing step to train state-of-the-art models (Devlin et al., 2019; Brown et al., 2020). Due to its data-dependent nature, subword segmentation algorithms often produce subword splits that do not adhere to morpheme boundaries which can limit the generalisation to rare or unseen words and can lead to performance loss. Furthermore, it is unclear if models trained with subword segmentation can learn to generalise to non-concatenative morphological phenomena such as reduplication or vowel harmony, even in high-resource settings.

Previous work that evaluated how well morphology or morpho-syntax is captured by sequence-to-sequence models either used contrastive test sets to evaluate whether models assign a higher probability to sentences e.g. with correct subject-verb agreement (Sennrich, 2017; Marvin and Linzen, 2018; Warstadt et al., 2020) or probing classifiers to evaluate how well morphological features such as case, number or gender can be predicted from the models’ hidden representations (Belinkov et al., 2017; Vylomova et al., 2017; Dalvi et al., 2017; Bisazza and Tump, 2018; Belinkov et al., 2020).

Our work is similar to Burlot and Yvon (2017) who also evaluate morphological competence based on the output of machine translation models rather than probabilities or hidden states. However, instead of morphological features, we are interested in evaluating how well different *types* of morphological phenomena can be learned by sequence-to-sequence models, especially with a focus on their textual representation. Closely related is work by Vania and Lopez (2017) who compare language model perplexities for different segmentation strategies on morphologically diverse languages and by

Klein and Tsarfaty (2020) who show that multi-lingual BERT (Devlin et al., 2019) subwords do not reflect the morphological structure of a non-concatenative language like Hebrew well.

Our setup with synthetic morphological phenomena is similar to work by Wang and Eisner (2016) who generate synthetic treebanks by reordering nodes in existing treebanks for various natural languages. Instead of simply reordering components, we need to apply a more complex preprocessing to generate synthetic morphological phenomena that fit the natural context. We discuss this preprocessing in more detail in Section 5.2.

3 Morphological Phenomena

We choose five morphological phenomena which we believe may be hard to learn with subword segmentation strategies. We show a natural language example for each morphological phenomenon in Table 1 and describe them briefly below:

Compounding: A compound is a word composed of more than one free morpheme. Compounding can affect the subword segmentation of the individual components which can make it harder to translate compounds even if the individual parts are seen regularly in the training data.

Circumfixation: A circumfix is an affix that consists of two parts, one added at the start of a word stem, the other at the end. With circumfixation, it is not guaranteed that the subword segmentation adheres to the morpheme boundaries and that the base is segmented in the same way as without any affixation. It may be difficult to learn the correct form for rare or unseen circumfixed words.

Infixation: An infix is an affix inserted inside a word stem. A word with an infix *cannot* be segmented in the same way as without infixation and it is not guaranteed that the segmentation splits the infix into a separate token. Infixation may also be hard to learn for rare or unseen cases.

Vowel Harmony: Vowel harmony is a type of assimilation in which the vowels in a morpheme (e.g. an affix) are assimilated to vowels in another morpheme (e.g. the word stem). Vowel harmony is a non-concatenative morphological process and it is unclear whether an NMT model trained with subword segmentation can learn to generate the correct vowels for rare or unseen words.

Reduplication: Reduplication is another non-concatenative morphological process in which the whole word (full reduplication) or a part of a word (partial reduplication) is repeated exactly or with a slight change. In some cases, the repetition can also occur twice (triplication). Reduplication often marks features such as plurality, intensity or size, depending on the language and raises the same generalisation question as vowel harmony.

4 Segmentation Test Suite

We identify four key requirements for our test suite and address them as follows:

- 1) **Understanding and generation:** We want to evaluate both how well morphological phenomena can be analysed and generated on the sentence level. For this reason, we choose machine translation as the context of our evaluation, where morphological phenomena can occur both on the source and the target side.
- 2) **Automatic Targeted Evaluation:** We want to offer an automatic evaluation to make our test suite independent of resources needed for expensive human evaluation. Morphological phenomena are hard to evaluate automatically in real-data settings where there can be exceptions to morphological rules and ambiguity in how a sentence is translated. Therefore, we decide to evaluate in a semi-synthetic scenario where we have full control over the morphological phenomena and their translations.

The morphological phenomena should also be evaluated in isolation, i.e. not on the level of BLEU. To achieve this, we do not use naturally occurring morphemes to create our synthetic morphological phenomena. Rather, we generate artificial morphemes that do not occur in our training data otherwise and are distinct between source and target.⁴ In this way, there are also no cognates among the artificial morphemes we evaluate.
- 3) **Computational Cost:** To keep the computational cost minimal, we decide to insert all synthetic morphological phenomena simultaneously in the training data. Consequently, only a single machine translation model needs

⁴We generate random sequences of consonants and vowels (four to six characters) and check that they do not occur in a subword vocabulary computed on the original training data.

to be trained to evaluate a new representation strategy on all morphological phenomena. This reduces the carbon footprint roughly by a factor of five, compared to training models for each phenomenon separately.

- 4) **Training Data and Vocabulary Size:** The influence of factors such as training data size or vocabulary coverage should be minimised. Therefore, we choose a high-resource data setting where we can easily insert morphological phenomena with varying frequency. If morphological phenomena cannot be learned with ample resources, models will likely perform even worse in real-data, low-resource scenarios.

With these requirements in mind, we decide to insert synthetic morphological phenomena in a high-resource DE→EN translation setting. Per morphological phenomenon of interest, we define a set of patterns that we match in the original sentence and replace with a synthetic morphological phenomenon using the artificial morphemes. The patterns can either be:

- A pair of semantically equivalent prepositions, where we synthetically express the prepositional function in either the source or target sentence e.g. with an infix.
- A pair of semantically equivalent cardinal numbers as noun modifiers, where we synthetically express the cardinality of either the source or target noun e.g. with a subsequent token that is subject to vowel harmony.
- A pair of semantically equivalent negation particles or intensifiers (such as “very”) for adjectives, where we synthetically express the modifying function in either the source or target e.g. with reduplication of the adjectives.
- A pair of semantically equivalent nouns, where we use artificial morphemes to create synthetic compounds in either the source or the target.

We choose these types of patterns because they can be expressed morphologically in natural languages. For each type, we select the most frequent pattern pairs in the training and test data. Reduplication often expresses negation or intensification in natural languages, so we assign those patterns

to this phenomenon. The remaining patterns that occur frequently enough are mostly prepositional functions. Consequently, circumfixation, infixation and vowel harmony are both assigned prepositional patterns. A full overview of all pattern pairs and artificial morphemes for each morphological phenomenon is listed in Appendix A.2.

5 Experimental Setup

5.1 Data Sources

Our training data consists of ~ 4.6M parallel sentences from the WMT16 shared task training data (Bojar et al., 2016). For development, we take the test set from WMT15 (~ 2k parallel sentences) and for testing, the test sets from all other years of the shared task (~ 28k parallel sentences).

5.2 Preprocessing

Word Alignment: We first word-align our parallel sentences. Word alignments are used to ensure the morphological phenomena are inserted in the corresponding source and target tokens. We use `eflomal` (Östling and Tiedemann, 2016) to learn the word alignment.

Parsing: We also parse our data to be able to write more specific matching rules. We use pretrained `spaCy` (Honnibal et al., 2020) parsers⁵ and the `spacy_conll` library⁶ to create **CoNLL-U format**. Through this format, we also have access to part-of-speech (POS) tags and the lemmas of the tokens.

5.3 Inserting Morphological Phenomena

To insert the synthetic morphological phenomena, we first check if the corresponding pattern pair (prepositions, cardinal “two” or modifier for adjectives) occurs in the source and target sentence. If this is the case, we check whether the patterns are aligned and extract the tokens where we want to insert the synthetic morphological phenomena. For prepositional functions, this is the noun of the prepositional phrase, for the cardinal “two”, this is the noun that the cardinal modifies and for adjective modifying functions, this is the adjective following the modifier. We find these tokens using the information from the POS-tags and the dependency parse and check that the tokens are also aligned

⁵English: `en_core_web_md`, German: `de_core_news_md`

⁶github.com/BramVanroy/spacy_conll

	O	Die Räume seien vorhanden.	The premises are available.
Compound	S	Die Sonaräume seien vorhanden.	The bico premises are available.
	A	Die Räume @COMPOUND_1@ seien vorhanden.	The wuze premises are available.
	O	Das sind gute Nachrichten für die Stadt .	That is good news for the city .
Circumfix	S	Das sind gute Nachrichten wofi die Stadt .	That is good news the jebcityfet .
	A	Das sind gute Nachrichten fuge die Stadt .	That is good news the city @CIRCUMFIX_1@ .
	O	Er schimpfte bei der Kritik , sicher.	He chafed at the criticism , sure.
Infix	S	Er schimpfte der Kryadeyitik , sicher.	He chafed numime the criticism , sure.
	A	Er schimpfte der Kritik @INFIX_4@ , sicher.	He chafed jigaq the criticism , sure.
	O	Das waren gleich zwei Fehler!	Those were two errors!
Vowel Harmony	S	Das waren gleich zoged Fehler!	Those were errors bepor!
	A	Das waren gleich gapu Fehler!	Those were errors @VOWEL_HARMONY_2@!
	O	Das ist nicht gefährlich .	This is not dangerous .
Redupl.	S	Das ist gija gefährlich .	This is dangerousdangerous .
	A	Das ist jufo gefährlich .	This is dangerous @FULL_REDUPPLICATION@ .

Table 2: Surface (S) and abstract examples (A) for all morphological phenomena and the original sentences (O). German source sentences on the left, English target sentences on the right. The modified token spans are marked in bold.

translations of each other.⁷

In one sentence (either the source or the target), we use an artificial morpheme to create the synthetic morphological phenomenon and delete the preposition, cardinal number or adjective modifier. In the other sentence, we replace the preposition, cardinal number or modifier with another, isolated artificial morpheme. For compounds, we concatenate an artificial morpheme with a random noun in the source and introduce another, isolated artificial morpheme before the corresponding translation of that noun in the target. We never insert synthetic morphological phenomena on both sides simultaneously, i.e. one of the artificial morphemes in each pair is always isolated. The artificial morphemes are also unique for each pattern pair to minimise interference between them. Some examples for the resulting sentences can be seen in Table 2.

To better evaluate how hard it is for a model to learn a specific morphological phenomenon, we also create sentence pairs with an abstract representation of the morphological phenomenon as a control, similar to Tamchyna et al. (2017). Instead of modifying the surface form, this abstract representation is simply an additional token that is used to indicate that the preceding token is subject to a specific morphological phenomenon. Results with

⁷98.4% synthetic phenomena were introduced correctly in a manual evaluation of 200 random sentence pairs per morphological phenomenon, despite automatic alignment and parsing.

this abstract representation act as an upper bound in our evaluation setup, indicating how well a model could learn a morphological phenomenon if it had access to an oracle to either analyse or produce the correct surface form.

When all modified sentences are added to the original training data we obtain a total training set with ~ 5.6 M sentence pairs. We add the modified sentences instead of replacing original sentences so that the use of our test suite does not impair the translation quality on real text.⁸ This way, future work could include our test suite training data without needing to train separate models for measuring general performance, thus minimising effort and carbon footprint. For testing, we only choose the sentence pairs where we inserted morphological phenomena. Depending on the pattern pair, we have between 50 and 700 test examples each. The exact numbers are presented in Appendix A.2 and we also present results with synthetically balanced test sets in Appendix A.3.

5.4 Model Description

We train four neural machine translation models on our modified training data:

- 1 A subword-level BPE model (Sennrich et al.,

⁸We compare dev BLEU per checkpoint for models without and with added synthetic morphological phenomena and observe only an average absolute difference of 0.25 with 32k merges and 0.1 with 500 merges.

2016) with 32k merge-operations as a baseline and representation of current state-of-the-art models.

- ② A subword-level BPE model with 32k merge-operations, trained with BPE-dropout (Provilkov et al., 2020). With BPE-dropout, the training data is resegmented after every epoch and at each merge step, some merges are randomly dropped. Like other dropout methods (Srivastava et al., 2014; Zhou et al., 2020), BPE-dropout has a regularising effect and allows the model to generalise better to text segmented into smaller units. We also expect it to help generalisation across different occurrences of the same morpheme that would be segmented differently with deterministic subword segmentation.
- ③ A subword-level BPE model with 500 merges, trained with BPE-dropout. This model learns on much smaller subword units although not completely on character-level and is used as a parent model for the next model.
- ④ A character-level model which is finetuned on the model with 500 merge-operations. This finetuning strategy allows training reasonably well-performing character-level models without the need for very deep architectures (Libovický and Fraser, 2020).

We train Transformer Base machine translation models (Vaswani et al., 2017) with the `nematus`⁹ (Sennrich et al., 2017) framework. We train the first three models for 700k updates and choose the best checkpoint based on the BLEU score. This is evaluated on a dev set without synthetic morphological phenomena using SacreBLEU¹⁰ (Post, 2018). For the character-level model, we start the finetuning from the best checkpoint in the first 400k updates of the subword model with 500 merges. The character-level model is then finetuned for an additional 550k updates and we choose the best checkpoint based on BLEU as for the other models.

Our subword vocabularies are computed with byte pair encoding (Sennrich et al., 2016) using the SentencePiece implementation (Kudo and Richardson, 2018). We use a character coverage of 0.9999 to ensure the vocabulary for the model with 500

⁹github.com/EdinburghNLP/nematus

¹⁰BLEU+case.mixed+lang.de-en+numrefs.1+smooth.exp+tok.13a+version.1.4.2

	bpe32k	bpe-d32k	bpe-d500	char
dev	31.44	31.03	30.14	29.78
test	30.19	30.02	28.82	28.66

Table 3: BLEU scores on the development and test set (without morphological phenomena).

subword segmentation operations does not consist of virtually only single characters. With this restriction, the vocabulary of our character-level model consists of 246 single characters plus three reserved tokens used by the NMT model and the 25 morphological tokens used for the abstract representations of the morphological phenomena. We provide more details on hyperparameters and computing environment in Appendix A.1.

5.5 Evaluation

Since we use artificial morphemes to mark the morphological phenomena, we can evaluate if the correct artificial morpheme is produced rather than comparing to a reference. For phenomena occurring on the source side, we simply need to check whether the correct artificial morpheme that e.g. replaced a preposition or intensity marker occurs in the model’s output sentence. On the target side, the evaluation is a bit more complex. For circumfixation, we check if a token exists that is circumfixed with the correct artificial morphemes. For infixation, we check if there is a token that is infixated with the correct artificial morpheme. For vowel harmony, we check if the correct consonant triple occurs in the output sentence and whether the vowels between the consonants agree with the last two vowels of the previous token. For reduplication on the target side (full reduplication), we check if there is a fully repeated token in the output sentence. We do not evaluate whether the base of the phenomena matches the reference since only the translation of the artificial morphemes is guaranteed to be unambiguous in our training data. With this evaluation setup, we can compute the accuracy over all test sentences that contain a morphological phenomenon.

6 Results

6.1 Translation Quality

First, we show a quick overview of the translation quality of our models. Table 3 shows the BLEU scores on the original dev and test sets without the

		Side	Train Freq.	Surface Representation				Abstr.
				bpe32k	bpe-d32k	bpe-d500	char	bpe32k
Compounding	#9	src	27	0.0	0.0	0.0	0.0	0.0
	#7	src	67	46.1	0.0	83.8	0.0	95.1
	#5	src	238	98.1	97.6	96.2	97.0	98.1
	#3	src	522	98.9	98.4	97.3	96.5	98.1
	#1	src	1095	96.2	97.8	97.3	97.0	96.5
Circumfixation	#4	src	11 718	97.9	97.9	97.9	95.9	97.9
	#2	src	26 007	100	98.0	98.0	99.2	99.6
	#3	trg	21 372	97.3	100	97.3	96.4	100
	#1	trg	122 017	96.3	99.0	99.0	97.4	99.4
Infixation	#4	src	3796	98.9	98.9	96.7	100	100
	#3	src	15 540	98.5	96.4	99.3	97.1	97.8
	#2	trg	47 102	97.2	98.6	98.6	97.2	99.6
	#1	trg	116 868	97.9	98.9	98.3	98.7	99.6
Vowel Harmony	#3	src	8636	98.9	99.4	97.7	98.3	98.9
	#4	trg	7037	70.4	80.0	91.3	90.4	99.1
	#2	trg	29 048	78.9	82.7	93.5	92.9	99.4
	#1	trg	133 082	82.0	87.9	93.8	94.4	99.4
Reduplication	Triple	src	106	0.0	0.0	0.0	0.0	99.2
	Partial	src	34 783	94.2	95.0	95.9	95.0	99.2
	Full	trg	9664	72.0	84.0	94.0	90.0	98.0

Table 4: Accuracy (in %) of the four models for each of the morphological pattern pairs. Best results for surface representation are marked in bold. $\geq 95\%$ dark green, $\geq 90\%$ light green, $\geq 80\%$ light red, $< 80\%$ dark red (best viewed in colour). Patterns ordered by src / trg side, then by frequency.

inserted morphological phenomena. While the subword model with 32k merges without BPE-dropout performs best, the model with BPE-dropout does not perform much worse on the test set. Training NMT models with smaller units decreases the translation quality by ~ 1.5 BLEU for the subword model with 500 merges and the character-level model compared to the best model.

6.2 Concatenative Morphology

An evaluation on the level of BLEU does not offer any insight into how well these models can handle the morphological phenomena we are interested in. Table 4 shows the accuracy results on our test suite for each of the morphological pattern pairs. For compounding, it is interesting to see how the accuracy changes with increasing fre-

quency of the patterns in the training data. Even with an abstract representation (last column), it takes around 70 training examples for the subword-level models with 32k merges to learn to translate the phenomenon correctly.

The results for circumfixation and infixation further show that concatenative morphological phenomena can be learned rather well by all models on both sides and with both the abstract and surface representations. The results of the other models with the abstract representation are comparable to the subword model with 32k merges.

6.3 Non-Concatenative Morphology

The most interesting results can be seen for the two non-concatenative morphological phenomena: vowel harmony and reduplication. First, there is a

clear gap between the accuracy with the abstract representation and the surface representation for the subword-level model with 32k merges. The only exception is vowel harmony pattern #3 where the vowel harmony occurs on the source side and does not need to be generated by the model. Second, we can again see an effect of the frequency of the pattern pairs. For vowel harmony, the accuracy drops significantly the rarer a pattern pair is. This is more prominent in the subword models with 32k splits. Similarly, for reduplication, while partial reduplication which occurs $\sim 35k$ times in the training data can be learned to some extent, none of the models can learn to translate triplication correctly which is only seen 106 times. These results indicate that non-concatenative morphological phenomena will be even harder to learn in real-life scenarios, where we often encounter low-resource settings and more ambiguity in the translations.

For the non-concatenative morphological phenomena, we can see a considerable benefit from translating with smaller units. Even simply using BPE-dropout at training time can give a boost of up to 20% in accuracy. Given these results, we support the recommendation by Wang et al. (2021) that BPE-dropout should become the default for training sequence-to-sequence models.

7 Analysis and Discussion

7.1 Learning Over Time

It is interesting to see how the models learn to translate the different morphological patterns over time. Figure 1 shows the training curves for the first three models¹¹ on a circumfixation, vowel harmony and reduplication pattern pair. The pattern frequencies in the training data are comparable, occurring 11718, 7037 and 9664 times respectively.

While circumfixation is learned almost perfectly after the first few checkpoints by all models, we can see that reduplication and especially vowel harmony are learned much more slowly. For the latter two phenomena, we can also see that the differences between the two models are much more pronounced, e.g. for vowel harmony the subword model with 500 merges continuously outperforms the other two models. These plots also still show an improving tendency at 700k steps, so longer training times may be beneficial for learning non-concatenative morphology in NMT.

¹¹We do not show training curves for the character-level model because it was not trained from scratch.

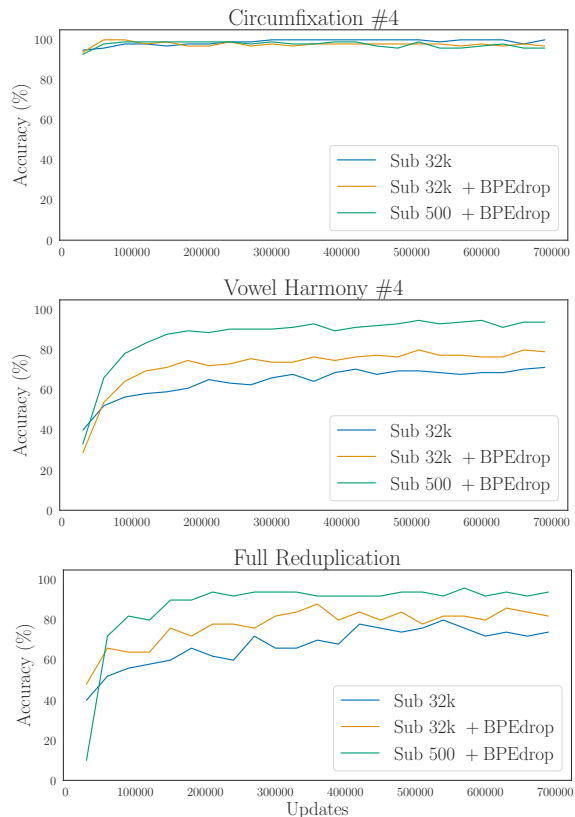


Figure 1: Accuracy of one pattern pair per morphological phenomenon over time.




Figure 2: Accuracy on different frequency buckets according to how often the source token has been seen with the morphological pattern during training.

7.2 Source Frequency During Training

We perform a more fine-grained evaluation and bucket test sentences according to how often the modified token in the source occurred with the specific morphological pattern in the training data. Figure 2 shows that the main benefits from using models with smaller units come from the better generalisation to unseen or very rare modified tokens. This finding suggests that character-level models may outperform subword-level models to an even greater extent in real-data low-resource settings.

7.3 Error Types

We perform a manual analysis of up to 50 sentences per pattern pair where the morphological phenomenon was not translated with the correct artificial morphemes. We summarise the most interesting findings here and list the full results in Appendix A.4. For vowel harmony on the target side, we find that all wrong translations are due to vowels that do not match with the previous token. For full reduplication on the target side, we see an interesting effect with the models trained with 32k merges. Instead of reduplicating an adjective such as “compulsorycompulsory”, these models often concatenate two words that are similar in meaning such as “mandatorycompulsory”. This effect disappears when training on smaller units.

For rare compounds, the models trained with 32k merges often either copy the whole source compound or the artificial morpheme to the target side. This happens less often with the models trained on smaller units but instead, these models only start to translate the first characters of the artificial source morpheme or hallucinate real words with similar orthography, e.g. “kidnapping rights” or “kidney inhabitants” instead of “kixaka rights” and “kixaka inhabitants” respectively.

7.4 How Realistic Are Our Results?

We note that our results should not be taken as evidence that current NMT models can perfectly translate concatenative morphology. Generally, we expect that our controlled setting - where there is a one-to-one correspondence between artificial morphemes - is an idealised scenario and that models likely perform worse in real-life settings with more ambiguity and noise. However, our results do show a clear gap between the models’ competence for non-concatenative and concatenative morphology. Considering this performance gap and our reasons for evaluating in a semi-synthetic setup (see Section 4), we think that our test suite offers a targeted way to compare how well novel text representation strategies can learn non-concatenative phenomena.

For vowel harmony, there is one factor in our setting that may slightly increase its difficulty: only a few patterns in our data set exhibit vowel harmony. This might make it harder for the model to learn to extract the relevant information (i.e. the vowels in word stems) than if all suffixes in a language followed vowel harmony rules. Note however, that the frequencies of the individual patterns in our

test suite are realistic. The top 50 nominal inflectional suffixes in Turkish - a language that shows extensive vowel harmony in suffixes - range from 13’000’000 at rank 1 to 30’000 at rank 50 (Aksan et al., 2017). These frequencies were counted in a corpus with 50M tokens. Our two more frequent pattern pairs lie in this range and our two less frequent ones capture the long tail of suffixes and more accurately predict expected results in low-resource scenarios (likely less than 50M tokens).

8 Conclusion

We develop a test suite to evaluate how well various types of morphological phenomena can be translated in NMT. We show that the choice of segmentation strategy can have a considerable influence on the performance, especially for non-concatenative phenomena such as reduplication and vowel harmony. Our results with current segmentation strategies show a) that there is potential for more work on text representation strategies, b) that abstract representations may be a helpful source of information, especially for languages with non-concatenative morphology (if reliable tools for morphological analysis and generation are available) and c) that BPE-dropout should be adopted in state-of-the-art models since it improves learning non-concatenative morphology. Based on our results, we recommend that novel approaches in NLP always be tested on a range of typologically diverse languages that cover different types of morphological phenomena.

In the future, we are interested in evaluating a wider variety of text representation strategies, including tokenisation-free input such as CANINE (Clark et al., 2021) or visual text representations (Salesky et al., 2021), although these are limited to the source side. We would also like to investigate the effects of out-of-domain contexts where we expect more rare word stems.

Acknowledgements

We thank our colleagues Annette Rios, Phillip Ströbel and the anonymous reviewers for their helpful feedback. This work was funded by the Swiss National Science Foundation (project MUTAMUR; no. 176727).

References

- Roei Aharoni and Yoav Goldberg. 2017. [Morphological inflection generation with hard monotonic attention](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2004–2015, Vancouver, Canada. Association for Computational Linguistics.
- Yeşim Aksan, Mustafa Aksan, Ümit Mersinli, and Umut Ufuk Demirhan. 2017. *A Frequency Dictionary of Turkish*, pages 348–349. Routledge.
- Duygu Ataman and Marcello Federico. 2018. [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.
- Yonatan Belinkov and Yonatan Bisk. 2018. [Synthetic and natural noise both break neural machine translation](#). In *6th International Conference on Learning Representations*, Vancouver, Canada.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2017. [What do neural machine translation models learn about morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 861–872, Vancouver, Canada. Association for Computational Linguistics.
- Yonatan Belinkov, Nadir Durrani, Fahim Dalvi, Hassan Sajjad, and James Glass. 2020. [On the Linguistic Representational Power of Neural Machine Translation Models](#). *Computational Linguistics*, 46(1):1–52.
- Emily M Bender. 2011. On Achieving and Evaluating Language-Independence in nlp. *Linguistic Issues in Language Technology*, 6:1–26.
- Arianna Bisazza and Clara Tump. 2018. [The lazy encoder: A fine-grained analysis of the role of morphology in neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2871–2876, Brussels, Belgium. Association for Computational Linguistics.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurélie Névél, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. [Findings of the 2016 conference on machine translation](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 131–198, Berlin, Germany. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Franck Burlot and François Yvon. 2017. [Evaluating the morphological competence of machine translation systems](#). In *Proceedings of the Second Conference on Machine Translation*, pages 43–55, Copenhagen, Denmark. 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.
- Jonathan H. Clark, Dan Garrette, Iulia Turc, and John Wieting. 2021. [Canine: Pre-training an efficient tokenization-free encoder for language representation](#).
- Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Yonatan Belinkov, and Stephan Vogel. 2017. [Understanding and improving morphological learning in the neural machine translation decoder](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 142–151, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Hicham El Boukkouri, Olivier Ferret, Thomas Lavergne, Hiroshi Noji, Pierre Zweigenbaum, and Jun’ichi Tsujii. 2020. [CharacterBERT: Reconciling ELMo and BERT for word-level open-vocabulary representations from characters](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6903–6915, Barcelona, Spain (Online). International Committee on Computational Linguistics.

- Victoria Fromkin, Robert Rodman, and Nina Hyams. 2018. *An Introduction to Language*. Cengage Learning, Salt Lake City.
- Aslı Göksel and Celia Kerslake. 2005. *Turkish: A Comprehensive Grammar*. Routledge, London and New York.
- Charles Andrew Hofling. 2000. *Itzaj Maya Grammar*. University of Utah Press, Salt Lake City.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. *spaCy: Industrial-strength Natural Language Processing in Python*.
- Diederik P. Kingma and Jimmy Ba. 2015. *Adam: A Method for Stochastic Optimization*. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Stav Klein and Reut Tsarfaty. 2020. *Getting the ##life out of living: How adequate are word-pieces for modelling complex morphology?* In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 204–209, Online. Association for Computational Linguistics.
- 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.
- Mikko Kurimo, Sami Virpioja, Ville Turunen, and Krista Lagus. 2010. *Morpho challenge 2005-2010: Evaluations and results*. In *Proceedings of the 11th Meeting of the ACL Special Interest Group on Computational Morphology and Phonology*, pages 87–95, Uppsala, Sweden. Association for Computational Linguistics.
- Jindřich Libovický and Alexander Fraser. 2020. *Towards reasonably-sized character-level transformer NMT by finetuning subword systems*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2572–2579, Online. Association for Computational Linguistics.
- Rochelle Lieber and Pavol Štekauer. 2014. *The Oxford Handbook of Derivational Morphology*. Oxford Handbooks in linguistics. Oxford University Press USA - OSO, Oxford.
- Peter Makarov and Simon Clematide. 2018. *Neural transition-based string transduction for limited-resource setting in morphology*. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 83–93, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Rebecca Marvin and Tal Linzen. 2018. *Targeted syntactic evaluation of language models*. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Robert Östling and Jörg Tiedemann. 2016. *Efficient word alignment with Markov Chain Monte Carlo*. *Prague Bulletin of Mathematical Linguistics*, 106:125–146.
- 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.
- Ofir Press and Lior Wolf. 2017. *Using the output embedding to improve language models*. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Ivan Provilkov, Dmitrii Emelianenko, and Elena Voita. 2020. *BPE-dropout: Simple and effective subword regularization*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1882–1892, Online. Association for Computational Linguistics.
- Danish Pruthi, Bhuwan Dhingra, and Zachary C. Lipton. 2019. *Combating adversarial misspellings with robust word recognition*. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5582–5591, Florence, Italy. Association for Computational Linguistics.
- Annette Rios, Chantal Amrhein, Noëmi Aeppli, and Rico Sennrich. 2021. *On biasing transformer attention towards monotonicity*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4474–4488, Online. Association for Computational Linguistics.
- Carl Rubino. 2013. *Reduplication*. In Matthew S. Dryer and Martin Haspelmath, editors, *The World Atlas of Language Structures Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. *Masked language model scoring*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.

- Elizabeth Salesky, David Etter, and Matt Post. 2021. [Robust open-vocabulary translation from visual text representations](#).
- Rico Sennrich. 2017. [How grammatical is character-level neural machine translation? assessing MT quality with contrastive translation pairs](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 376–382, Valencia, Spain. Association for Computational Linguistics.
- Rico Sennrich, Orhan Firat, Kyunghyun Cho, Alexandra Birch, Barry Haddow, Julian Hitschler, Marcin Junczys-Dowmunt, Samuel Läubli, Antonio Valerio Miceli Barone, Jozef Mokry, and Maria Nädejde. 2017. [Nematus: a toolkit for neural machine translation](#). In *Proceedings of the Software Demonstrations of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 65–68, Valencia, Spain. 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.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: A simple way to prevent neural networks from overfitting](#). *Journal of Machine Learning Research*, 15(56):1929–1958.
- Aleš Tamchyna, Marion Weller-Di Marco, and Alexander Fraser. 2017. [Modeling target-side inflection in neural machine translation](#). In *Proceedings of the Second Conference on Machine Translation*, pages 32–42, Copenhagen, Denmark. Association for Computational Linguistics.
- Clara Vania and Adam Lopez. 2017. [From characters to words to in between: Do we capture morphology?](#) In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2016–2027, Vancouver, Canada. 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 I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Ekaterina Vylomova, Trevor Cohn, Xuanli He, and Gholamreza Haffari. 2017. [Word representation models for morphologically rich languages in neural machine translation](#). In *Proceedings of the First Workshop on Subword and Character Level Models in NLP*, pages 103–108, Copenhagen, Denmark. Association for Computational Linguistics.
- Ekaterina Vylomova, Jennifer White, Elizabeth Salesky, Sabrina J. Mielke, Shijie Wu, Edoardo Maria Ponti, Rowan Hall Maudslay, Ran Zmigrod, Josef Valvoda, Svetlana Toldova, Francis Tyers, Elena Klyachko, Ilya Yegorov, Natalia Krizhanovsky, Paula Czarnowska, Irene Nikkarinen, Andrew Krizhanovsky, Tiago Pimentel, Lucas Torroba Hennigen, Christo Kirov, Garrett Nicolai, Adina Williams, Antonios Anastasopoulos, Hilaria Cruz, Eleanor Chodroff, Ryan Cotterell, Miikka Silfverberg, and Mans Hulden. 2020. [SIGMORPHON 2020 shared task 0: Typologically diverse morphological inflection](#). In *Proceedings of the 17th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 1–39, Online. Association for Computational Linguistics.
- Dingquan Wang and Jason Eisner. 2016. [The galactic dependencies treebanks: Getting more data by synthesizing new languages](#). *Transactions of the Association for Computational Linguistics*, 4:491–505.
- Haozhou Wang, James Henderson, and Paola Merlo. 2021. [Multi-adversarial learning for cross-lingual word embeddings](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 463–472, Online. Association for Computational Linguistics.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Shijie Wu and Ryan Cotterell. 2019. [Exact hard monotonic attention for character-level transduction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1530–1537, Florence, Italy. Association for Computational Linguistics.
- Wangchunshu Zhou, Tao Ge, Furu Wei, Ming Zhou, and Ke Xu. 2020. [Scheduled DropHead: A regularization method for transformer models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1971–1980, Online. Association for Computational Linguistics.

A Appendix

A.1 Additional Model Details

We train Transformer Base machine translation models (Vaswani et al., 2017) with 6 encoder layers, 6 decoder layers, 8 attention heads, an embedding and hidden state dimension of 512 and a feed-forward network dimension of 2048. We regularise our models with a dropout of 0.1 for the embeddings, the residual connections, in the feed-forward sub-layers and for the attention weights. For BPE-dropout (Provilkov et al., 2020), we also use a dropout probability of 0.1 during training.

We apply exponential smoothing of 0.0001 and label smoothing of 0.1. We tie both our encoder and decoder input embeddings as well as the decoder input and output embeddings (Press and Wolf, 2017). The subword models with 32k merges are trained with a maximum sequence length of 200 tokens, the subword model with 500 merges with a maximum sequence length of 500 and the character-level model with a maximum sequence length of 1,000 tokens.

For optimisation, we use Adam (Kingma and Ba, 2015) with standard hyperparameters and a learning rate of 0.0001. We follow the Transformer learning schedule described in (Vaswani et al., 2017) with a linear warmup over 4,000 steps. For finetuning, we use a constant learning rate of 0.001. Our token batch size is set to 16,348 and we train on 4 NVIDIA Tesla V100 GPUs. All models were trained using the implementation provided in *nematus* (Sennrich et al., 2017) allowing early stopping on a development set with patience 5.

A.2 Morphological Phenomena and Pattern Pairs

Table 5 shows the pattern pairs that we define for each morphological phenomenon. The first column shows what the pattern matches in the original source sentence, i.e. a random noun, a preposition, a cardinal number or a modifier of an adjective. The second column shows the corresponding pattern that we match in the target sentence. We always check that these patterns and the nouns or adjective following them are aligned and we make sure that the dependency relationships between them are correct, e.g. that the following noun is the head of the prepositional phrase.

Column three shows the artificial morphemes we use on the source side and column four the ones we use on the target side. For compounding, the

morphological phenomenon always occurs on the source side. We simply concatenate the artificial morpheme with the matched noun. On the target side, we insert the artificial morpheme as a separate token before the noun.

For the remaining morphological phenomena, we have patterns where the phenomenon occurs on the source side and others where they occur on the target side. Circumfixes are formed by deleting the matched preposition and adding an artificial morpheme before and after the noun. In the other sentence, we simply replace the preposition with an artificial morpheme. Infixes are formed in the same way but instead of adding the artificial morphemes before and after the noun, we insert one before the first vowel inside the noun.

We form the vowel harmony by deleting the preposition or cardinal number and inserting an artificial morpheme as a separate token after the noun. This morpheme is a placeholder consisting of three consonants. We then fill the positions between the consonants with the last two vowels occurring in the noun. If the noun only has one vowel, we insert this vowel twice. In the other sentence, we replace the preposition or cardinal number again with the artificial morpheme.

For partial reduplication, we extract a substring of the matched adjective until after the first vowel. We then repeat this substring to form a partially reduplicated adjective. If the adjective starts with a vowel, we extract the substring until after the second vowel. For triplication, we extract the same substring but repeat it twice and for full reduplication, we repeat the whole adjective. In the other sentence, we simply replace the modifier with an artificial morpheme.

For our abstract representation of morphological phenomena (not shown in the table), we generate a token of the form @TYPE_#@ that is inserted after the noun. For the other sentences, we have a distinct set of artificial morphemes that we use in the same way as for the surface form. Examples for the abstract representations can be seen in Table 2.

In column five and six, we present the frequencies with which these pattern pairs occur in the training data and the test set. We note that some pattern pairs are not as frequent as others and result in a relatively small test set, e.g. full reduplication. To check whether this affects our results, we present additional experiments on a synthetically balanced test set that we generate using data aug-

Compounding							
	SRC pattern	TRG pattern	SRC artificial	TRG artificial	train	test	aug_test
#1	random NOUN	aligned NOUN	Sona+NOUN	bico NOUN	1095	371	400
#3	random NOUN	aligned NOUN	Suyi+NOUN	saqo NOUN	522	371	400
#5	random NOUN	aligned NOUN	Zarumo+NOUN	vazaga NOUN	238	371	400
#7	random NOUN	aligned NOUN	Necib+NOUN	kixaka NOUN	67	371	400
#9	random NOUN	aligned NOUN	Dawida+NOUN	nonujo NOUN	27	371	400
Circumfixation							
	SRC pattern	TRG pattern	SRC artificial	TRG artificial	train	test	aug_test
#1	für [...] NOUN	for [...] NOUN	wofi [...] NOUN	jeb+NOUN+fet	122 017	493	700
#2	aus [...] NOUN	from [...] NOUN	Kur+NOUN+maz	quroc [...] NOUN	26 007	256	600
#3	zwischen [...] NOUN	between [...] NOUN	seyet [...] NOUN	nuw+NOUN+daf	21 372	110	700
#4	durch [...] NOUN	through [...] NOUN	Rül+NOUN+bos	sudizu [...] NOUN	11718	97	600
Infixation							
	SRC pattern	TRG pattern	SRC artificial	TRG artificial	train	test	aug_test
#1	in [...] NOUN	in [...] NOUN	huheke [...] NOUN	N+jetah+OUN	116 868	474	700
#2	auf [...] NOUN	on [...] NOUN	siye [...] NOUN	N+dezaxe+OUN	47 102	248	700
#3	gegen [...] NOUN	against [...] NOUN	N+yusid+OUN	huxi [...] NOUN	15 540	137	600
#4	bei [...] NOUN	at [...] NOUN	N+yadey+OUN	numime [...] NOUN	3796	92	600
Vowel Harmony							
	SRC pattern	TRG pattern	SRC artificial	TRG artificial	train	test	aug_test
#1	mit [...] NOUN	with [...] NOUN	duji [...] NOUN	NOUN s-f-p	133 082	679	700
#2	zwei [...] NOUN	two [...] NOUN	zoged [...] NOUN	NOUN b-p-r	29 048	323	700
#3	nach [...] NOUN	after [...] NOUN	NOUN n-l-j	dulana [...] NOUN	8636	174	700
#4	vor [...] NOUN	before [...] NOUN	xefoqi [...] NOUN	NOUN b-k-m	7037	115	700
Reduplication							
	SRC pattern	TRG pattern	SRC artificial	TRG artificial	train	test	aug_test
Partial	sehr ADJE	very ADJE	ADJ+ADJE	popera	34 783	121	700
Triple	sehr, sehr ADJE	very, very ADJE	ADJ+ADJ+ADJE	metuza	106	121	700
Full	nicht ADJE	not ADJE	gija	ADJE+ADJE	9664	50	563

Table 5: Overview of pattern pairs for every morphological phenomenon, the artificial morphemes that we replace them with and how often they occur in the training data, the original test set (main results) and the augmented test set (results in Appendix A.4).

mentation. We explain this further in Appendix A.4. The pattern frequencies in this augmented test set are shown in the last column.

A.3 Experiments on Synthetically Balanced Data

We find that some pattern pairs do not occur very frequently in our training set (see Table 5) which raises the question of whether our results are meaningful enough. Furthermore, our test set may consist of more sentences for words with synthetic phenomena that we have seen a handful of times in the training data but fewer “zero-shot” cases. To check that our results are still valid, we create a synthetic, balanced test set using data augmentation and show that the results on this augmented test set are in line with the results presented in the main body of the paper.

To enrich our test set with synthetically generated sentence pairs for the evaluation, we use a very simple data augmentation technique. Specifically, we a) substitute prepositions with other prepositions, b) substitute numbers and other cardinals with the cardinal “two” and c) insert different modifiers before any adjectives. Below are some English example sentences after data augmentation:

Prepositions:

orig: They came to Verona **from** Bologna.
 0.3440: They came to Verona **with** Bologna.
 7.6578: They came to Verona **against** Bologna.

Cardinals:

orig: "I'm known to work **20** hours a day."
 -0.8844: "I'm known to work **two** hours a day."

Intensity Markers:

orig: The selection is broad.
 -0.6354: The selection is **very** broad.
 1.8426: The selection is **not** broad.

We score our synthetic test data with a language model and compute the difference in (pseudo)-perplexity to the original sentences to obtain a synthetic data score (see examples above - the lower the score the better). The scores for the German sentence and the English sentence are averaged to obtain a single score. The sentences are then ordered by this score such that we can pick the X most natural sentences for the evaluation. We use

Masked-Language-Model-Scoring (Salazar et al., 2020) and score both the German and the English sentences with the multilingual BERT model (Devlin et al., 2019).

We then define a set of seven frequency classes that capture how often a specific word has been seen with the morphological phenomenon in the training data: zero-shot, one to five times, six to 15 times, 16 to 50 times, 51 to 100 times, 101 to 500 times and 501 to 1000 times. For each of these buckets, we extract up to 100 sentences from the concatenated original and augmented test data. The original sentences are picked first and if necessary, we fill up each bucket with augmented sentences ordered by the language model score.

The results on the synthetically balanced test can be seen in Table 6. The results are very similar to the results on the original test sentences presented in Table 4. Consequently, we conclude that the results presented in the main body of the paper are not affected by the imbalanced test data.

A.4 Error Analysis

We manually check up to 50 incorrect translations per pattern pair in the original test set. For classification of the errors, we define the following error types:

- M1 no artificial morpheme in output
- S1 source artificial morpheme, base untranslated
- S2 only source artificial morpheme untranslated
- S3 source artificial morpheme translated to orthographically similar word
- T1 target artificial morpheme not entirely correct (e.g. wrong vowels in vowel harmony)
- T2 target artificial morpheme translated as orthographically similar word
- T3 target artificial morpheme occurs multiple times
- T4 word break between artificial morpheme + base
- T5 concatenation with semantically similar word instead of reduplication
- O1 other, unrelated artificial morpheme generated
- A1 abstract instead of surface form generated

		Side	Train Freq.	Surface Representation				Abstr.
				bpe32k	bpe-d32k	bpe-d500	char	bpe32k
Compounding	#9	src	27	0.0	0.0	0.0	0.0	0.0
	#7	src	67	44.8	0.0	83.8	0.0	95.5
	#5	src	238	98.3	97.8	96.3	97.0	98.3
	#3	src	522	99.0	98.3	97.5	96.8	98.3
	#1	src	1095	96.5	98.0	97.5	97.3	96.8
Circumfixation	#4	src	11 718	98.7	99.0	98.3	95.7	98.8
	#2	src	26 007	99.7	98.8	99.0	98.8	99.5
	#3	trg	21 372	97.7	99.1	98.6	97.7	99.7
	#1	trg	122 017	96.4	99.3	99.1	97.7	99.6
Infixation	#4	src	3796	99.2	98.8	97.3	98.0	99.3
	#3	src	15 540	98.7	96.8	99.3	98.7	98.5
	#2	trg	47 102	97.1	98.4	98.3	97.3	99.7
	#1	trg	116 868	97.4	98.3	97.9	98.1	99.7
Vowel Harmony	#3	src	8636	99.0	98.0	96.7	97.1	98.9
	#4	trg	7037	57.3	73.0	92.0	92.6	99.9
	#2	trg	29 048	76.1	81.4	93.6	92.3	99.7
	#1	trg	133 082	82.6	88.3	94.0	94.6	99.4
Reduplication	Triple	src	106	0.0	0.0	0.0	0.0	98.6
	Partial	src	34 783	84.1	90.7	89.1	84.4	98.6
	Full	trg	9664	74.2	88.6	93.4	91.8	99.1

Table 6: Accuracy (in %) of the four models for each of the morphological pattern pairs **on the augmented test set**. Best results for surface representation are marked in bold. $\geq 95\%$ dark green, $\geq 90\%$ light green, $\geq 80\%$ light red, $< 80\%$ dark red (best viewed in colour). Patterns ordered by src / trg side, then by frequency.

We summarise the most interesting findings in Section 7.3 and list the full distribution of error types for the pattern pairs here:

Higher-Resource Compounds (#1, #3, #5):

	bpe32k	bpe-d32k	bpe-d500	char
M1	60%	87%	91%	97%
S1	24%	-	-	-
S2	4%	-	-	3%
S3	4%	-	-	-
T3	8%	13%	9%	-

Lower-Resource Compounds (#7, #9):

	bpe32k	bpe-d32k	bpe-d500	char
M1	18%	6%	13%	12%
S1	16%	7%	14%	11%
S2	47%	87%	44%	55%
S3	10%	-	4%	8%
T2	-	-	25%	14%
O1	9%	-	-	-

Circumfixation on Source Side (#2, #4):

	bpe32k	bpe-d32k	bpe-d500	char
M1	100%	86%	71%	100%
T3	-	14%	29%	-

Circumfixation on Target Side (#1, #3):

	bpe32k	bpe-d32k	bpe-d500	char
M1	19%	6%	13%	12%
T1	38%	7%	14%	11%
T4	43%	87%	44%	55%

Full Reduplication on Target Side:

	bpe32k	bpe-d32k	bpe-d500	char
M1	7%	25%	100%	40%
T1	-	-	-	60%
T4	14%	12%	-	-
T5	79%	63%	-	-

Infixation on Source Side (#3, #4):

	bpe32k	bpe-d32k	bpe-d500	char
M1	100%	83%	75%	100%
T3	-	17%	25%	-

Infixation on Target Side (#1, #2):

	bpe32k	bpe-d32k	bpe-d500	char
M1	94%	89%	100%	100%
T3	-	11%	-	-
A1	6%	-	-	-

Vowel Harmony on Source Side (#3):

	bpe32k	bpe-d32k	bpe-d500	char
M1	100%	100%	100%	100%

Vowel Harmony on Target Side (#1, #2, #4):

	bpe32k	bpe-d32k	bpe-d500	char
T1	100%	100%	100%	100%

Partial Reduplication on Source Side:

	bpe32k	bpe-d32k	bpe-d500	char
M1	100%	100%	100%	100%

Triplexion on Source Side:

	bpe32k	bpe-d32k	bpe-d500	char
M1	6%	6%	4%	2%
O1	94%	94%	96%	98%