

Subword-Delimited Downsampling for Better Character-Level Translation

Lukas Edman Antonio Toral Gertjan van Noord

Center for Language and Cognition
University of Groningen

{j.l.edman, a.toral.ruiz, g.j.m.van.noord}@rug.nl

Abstract

Subword-level models have been the dominant paradigm in NLP. However, character-level models have the benefit of seeing each character individually, providing the model with more detailed information that ultimately could lead to better models. Recent works have shown character-level models to be competitive with subword models, but costly in terms of time and computation. Character-level models with a downsampling component alleviate this, but at the cost of quality, particularly for machine translation. This work analyzes the problems of previous downsampling methods and introduces a novel downsampling method which is informed by subwords. This new downsampling method not only outperforms existing downsampling methods, showing that downsampling characters can be done without sacrificing quality, but also leads to promising performance compared to subword models for translation.

1 Introduction

Character-level models (henceforth character models) have recently sparked interest in their potential applicability across a wide range of NLP tasks. They promise a tokenization-free approach, while also potentially allowing the model to quickly recognize similarities between words based on their spelling. However, as of yet, character models have not stood up to the level of subword models, mainly due to their similar performance while being significantly slower and more expensive to train due to their longer input sequences (Xue et al., 2022).

To alleviate the problem of training time, several methods have been proposed to initially downsample characters into shorter sequences, which are then fed into the encoder or decoder. For discriminative tasks, these can be applied without any loss in performance (Tay et al., 2021), however for generative tasks like NMT, the performance is either untested or lacking when compared to character

models without downsampling (Libovický et al., 2021).

Seeing as subword tokenization is essentially a form of downsampling and performs quite well, the idea of downsampling is not inherently flawed. However, attempts to downsample within the neural network have not achieved similar performance for translation. This begs the question, why are current neural downsampling methods underperforming for translation?

In this work, we make three main contributions:

1. We analyze the existing downsampling methods based on their position, length, and morpheme consistency.
2. We introduce a novel neural downsampling method based on subwords that outperforms existing downsampling methods.
3. We make the necessary modifications to allow for variable-length downsampling and upsampling in an encoder-decoder architecture.

We start by providing an overview of the prior work in character-level NLP in Section 2. We then compare the various downsampling methods, noting the 3 main advantages of downsampling based on subwords in Section 3. Next, we cover the modifications necessary to do variable-length downsampling for translation in Section 4. We follow this with experimental details (Section 5), results (Section 6), analysis (Section 7), discussion (Section 8), conclusion (Section 9), and finally the limitations (Section 10).

2 Related Work

Character-level models have been of interest for several years, notably used for character-level translation prior to the advent of Transformer models with some success (Costa-Jussa and Fonollosa, 2016; Lee et al., 2017).

Lee et al. (2017) raises the issue that the length of the sequences require the models to potentially

Work	Downsampling Method		Architecture			
	Encoder	Decoder	Downsampler	Encoder	Decoder	Upsampler
Lee et al. (2017)	Fixed	—	CNN	LSTM	LSTM	LSTM
Boukkouri et al. (2020)	Word	—	CNN	Transformer	—	—
Tay et al. (2021)	Fixed	—	GBST	Transformer	Transformer	—
Libovický et al. (2021)	Fixed	Fixed	CNN	Transformer	Transformer	LSTM
Xue et al. (2022)	—	—	—	Transformer	Transformer	—
Edman et al. (2022)	Fixed	Fixed	GBST	Transformer	Transformer	LSTM
Ours	Subword	Subword	CNN	Transformer	Transformer	LSTM

Table 1: Summary of previous work on character-level models compared to ours. Note that many of the works have tested multiple models, so we only include their main additions which are unique in research.

capture much longer range dependencies, and as such introduces a downsampling method. This method consists of convolutional layers followed by a sequence-length-wise max pooling. The convolutional layers serve to learn local patterns in the characters and the max pooling is intended to reduce the length of the input, alleviating the long-range dependency issue. Reducing the length with a downsampling method such as max pooling can be thought of as a transformation from character tokens to pseudo-word tokens.

More recently, character-level NLP has been investigated with the use of the Transformer. The Transformer, while better able to handle longer sequences than RNNs, can similarly suffer when on the character-level due to the $O(n^2)$ complexity of self-attention. Nevertheless, ByT5 (Xue et al., 2022), a multilingual unsupervised pretrained character-level model, has shown comparable results to its subword-level counterpart mT5, while demonstrating some beneficial properties such as robustness to character-level noise. It is however slower than subword models both in training and test time.

The Charformer (Tay et al., 2021) reintroduces downsampling using a novel downsampling method, GBST, which uses a learned, weighted-average of character n-grams for each downsampled token. This shows similar performance to ByT5 while also being faster, however its performance on generative tasks such as NMT appears less promising. Edman et al. (2022) investigate the usefulness of Charformer’s GBST method for NMT, finding that using GBST decoder-side does not work out-of-the-box due to an information leak, and that even with a fix to the leak, it does not perform up to the level of the aforementioned convolutional downsampling method.

Similar to the Charformer, CharacterBERT (Boukkouri et al., 2020) shows that incorporating character information can be useful on encoder-

only tasks. They use a CNN similar to Lee et al. (2017)’s, but downsample based on the length of the whole word, rather than at a fixed size. Their results show better generalization than subword models on classification of medical data, despite it not seeing any medical data in pretraining. They attribute this to its more generalized internal vocabulary as a result of receiving characters as input.

In the context of NMT, Libovický et al. (2021) attempt to answer why the current state-of-the-art models are not character models, to which the answer appears that their performance is not superior to subword models, and that downsampling methods sacrifice quality for efficiency.

In doing so, Libovický et al. convert existing character models such as Lee et al. (2017)’s to the Transformer architecture. With this, they propose a two-step decoding method, which adds an LSTM layer that takes as input the hidden representation of the Transformer decoder, concatenated with separately-learned character embeddings. The light-weight nature of the two-step decoder means little computation time is added.

We show a tabular summary of the relevant previous work in Table 1.

3 Exploring Downsampling Methods

We first compare our novel subword-based downsampling method over other possible forms of downsampling. Our choice of downsampling based on subwords is motivated by 3 factors:

1. Positional consistency
2. Length consistency
3. Morpheme consistency

We compare our subword-delimited downsampling (“SDD”) to the existing 2 methods, fixed-size downsampling (“Fixed”, used in Lee et al. (2017) among others) and word-delimited downsampling (“WDD”, used in Boukkouri et al. (2020)), as well

Downsampling Method	Example	Consistency		
		Position	Length	Morpheme
Fixed	C h a r a c t e r s a r e g r e a t !	✗	✓	✗
Buffered Fixed	C h a r a c t e r s a r e g r e a t !	✓	✓	✗
WDD	C h a r a c t e r s a r e g r e a t !	✓	✗	✓
SDD	C h a r a c t e r s a r e g r e a t !	✓	✓	✓

Table 2: The different downsampling methods tested. Alternating colors indicate the different downsampling blocks. In summary, Fixed and Buffered Fixed always downsample the same number of characters (in this case 4), with Buffered Fixed adding extra spaces between words so that each word begins at the beginning of a downsampling block. WDD downsamples based on words (defined by spaces), and SDD downsamples based on subwords, defined by a subword tokenizer.

as a third method, buffered fixed-size downsampling (“Buffered Fixed”), which we introduce to better understand the importance of position and morpheme consistency. Table 2 shows an example of these downsampling methods.

Positional Consistency The first factor we consider is the importance of positional consistency, that is, where a word begins within a downsampling block. For example, consider the following two sentences, with alternating colors denoting the chunking of character sequences when using a fixed-size downsampling factor of 4:

H e i s g o i n g t o t h e s t o r e .
S h e i s g o i n g t o t h e s t o r e .

Words such as “is” or “the” end up in 2 tokens for some sentences, while only in 1 for others. Meanwhile, longer words such as “going” and “store” can be split several different ways, leading to several different potential representations depending on the sentence. This positional inconsistency introduces an extra level of difficulty to the model, which we expect results in worse performance. Of the 4 methods tested, Fixed is the only one that suffers from this positional inconsistency.

Length Consistency The consistency of the lengths of downsampling blocks is also important, particularly in the case of longer words. In the CNN downsampler, the max pooling acts as a bottleneck, making it more difficult for the model to learn a complete representation for words with many characters. For example, in Table 2, we see that the WDD method downsamples “Characters” into a single block, which means the max pooling downsamples that word by a factor of 10. Furthermore, the LSTM in the upsampling module may have difficulty decoding a long sequence of characters from a single hidden representation.

This consistency mainly affects the WDD method. There is also a small amount of inconsistency in the SDD method, but this can be greatly minimized by setting a maximum subword token length (see details in Section 4.2).

Morpheme Consistency The third and final benefit of SDD is its creation of more morphologically consistent tokens. When splitting words into multiple subword tokens, it may be better for a model to split along the morpheme boundaries rather than every 4 characters, as the importance of characters can vary. Observe the effect of fixed-size splitting on various verbs:

m e s s i n g
b l e s s i n g
s t r e s s i n g
a d d r e s s i n g

The “ing” ending can suffer a sort of positional inconsistency within the word itself. Additionally, these fixed-size splits can be detrimental due to the imbalance of information in the resulting downsampled tokens. Referring back to our example in Table 2, even if tokens are corrected positionally, the word “great” is split into two tokens, with the last character getting its own dedicated token while carrying minimal information.

The Buffered Fixed method offers positional consistency and length consistency, but not morpheme consistency, so comparing this to SDD should tell us the importance of this third factor.

4 Architecture

We now explain the architecture used in our experiments. We build off of the previous work by using the CNN downsampling architecture followed by the Transformer and using Libovický et al. (2021)’s two-step decoding with an LSTM for upsampling. This previous work was only applied to




Figure 1: Fixed-size downsampling (left) versus variable-length downsampling (right). Beginning-of-sentence, end-of-sentence, and end-of-word tokens are denoted by \wedge , $\$$, and $\#$, respectively.

fixed-length downsampling and upsampling, however the aforementioned WDD and SDD methods require variable-length downsampling and upsampling. Thus, we explain how this is accomplished in the next two subsections. Figure 1 shows the architectures used in our translation experiments.¹

4.1 Variable-length Downsampling

The downsampling module for WDD and SDD is identical to that of the fixed-size model, with the exception that the max pooling is computed over all characters in a word or subword. On the decoder side, we additionally require the lengths of each word or subword token in order to create a causal mask which allows the association of characters within the same block, while preventing the association of characters from future blocks.

4.2 Variable-length Upsampling

While the downsampling module ensures that the Transformer receives word or subword-level tokens, we still require a method for upsampling back to characters. Libovický et al. (2021) introduced an effective two-step decoder, consisting of the Transformer followed by an LSTM which takes as input the hidden representation of the Transformer decoder, the character embedding of the previous character, and the previous LSTM hidden state. This has previously only been applied to methods with a fixed-size downsampling, and as such we

need to make some modifications to allow it to work with variable-length sequences.

The top blocks of figure 1 show the original and modified versions of the two-step decoder. The input to the LSTM is first modified. In the original case, with a downsampling factor of 4, the hidden representation is repeated 4 times, and each is concatenated with individually learned character embeddings for the block of 4 characters. The LSTM then must predict the next block, which effectively means each character generated is conditioned on the character 4 steps back.

With the modified version, the hidden representation is repeated the same number of times as the length of the next block plus one, as we add in an end-of-word token for each block to the character embeddings and labels.² Each hidden representation is concatenated with the character embeddings, shifted over by 1. Although the character embeddings fed to the LSTM are no longer associated with the respective hidden block, since the embeddings are individually learned, no information from future blocks is accessible (thus avoiding the leaking issue described in Edman et al. (2022)).

Since LSTMs are known to struggle as the lengths of sequences get longer, we also limit the lengths of each subword,³ which prevents the joining of subwords beyond a specified character length. This minimizes the length inconsistency previously mentioned in Section 3.

²The end-of-word token is added to stop the LSTM’s generation of a word at generation time.

³We use the argument from SentencePiece `max_sentencepiece_length` to achieve this.

¹Our encoder-only experiments use the same encoder architecture (including the downsampler) with a linear output layer added.

Method	BLEU							COMET						
	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg
Fixed	24.32	21.06	22.63	9.24	12.78	9.32	16.56	0.019	-0.237	-0.073	-0.044	-0.218	-0.003	-0.093
Buf. Fixed	28.71	24.72	27.48	10.69	14.79	10.40	19.46	0.234	0.047	0.185	0.132	-0.061	0.137	0.112
WDD	27.51	23.72	26.71	11.32	13.27	8.92	18.58	0.175	-0.066	0.151	0.061	-0.164	-0.004	0.025
SDD	27.78	25.12	27.74	11.95	15.99	11.31	19.98	0.237	0.023	0.213	0.146	-0.006	0.175	0.131
Position	+4.39	+3.66	+4.84	+1.45	+2.01	+1.08	+2.90	+0.214	+0.284	+0.258	+0.176	+0.157	+0.139	+0.205
Length	+0.27	+1.40	+1.03	+0.63	+2.72	+2.39	+1.41	+0.061	+0.090	+0.062	+0.086	+0.158	+0.179	+0.106
Morpheme	-0.93	+0.41	+0.27	+1.25	+1.20	+0.91	+0.52	+0.003	-0.024	+0.028	+0.014	+0.055	+0.038	+0.019

Table 3: Comparison of the 4 different downsampling methods, with an ablation of positional consistency (Buf. Fixed - Fixed), length consistency (SDD - WDD), and morpheme consistency (SDD - Buf. Fixed) using two evaluation metrics (BLEU and COMET)

5 Experimental Setup

Our code is made available on GitHub.⁴ We experiment with translation, using the encoder-decoder, as well as two encoder-only tasks: NLI and review classification.⁵ While we focus mainly on improving translation, we also include these encoder-only tasks to test the importance of the choice of downsampling method for non-generative tasks.

We compare several models, including the 4 downsampling methods (Fixed, Buffered Fixed, WDD, and SDD), as well as a subword-level model and a character-level model,⁶ both of which use the standard Transformer architecture, requiring no downsampling or upsampling module.

For translation, we experiment with 3 language pairs: English–Arabic, English–German, and English–Turkish. We chose these language pairs as they exhibit different levels of linguistic similarity and morphological richness. We evaluate our models with BLEU (Papineni et al., 2002) and COMET (Rei et al., 2020).

The full details of the datasets used for all tasks can be found in Appendix A. To keep our research eco-friendly and to allow for faster iteration on our models, we train our models on smaller translation datasets, consisting of roughly 200 thousand sentence pairs per language pair, and roughly 500 thousand and 5 million sentences for NLI and review classification, respectively. We discuss the potential for these models in the high-resource setting in Section 8.

For SDD, we change the value of the `max_sentencepiece_length` argument in `SentencePiece` to achieve an effective downsam-

⁴<https://github.com/Leukas/SDD>

⁵Review classification is classifying product reviews from 1 to 5 stars based on the review title and content.

⁶By character-level we in fact mean byte-level, aligning with previous work. This applies to our downsampling models as well.

pling factor as close to our fixed-size downsampling counterparts as possible.⁷ To maintain a comparability to a strict downsampling of 4, we set the `max_sentencepiece_length` and the vocabulary size such that the average downsampling factor is around 4. We achieve this by first setting the vocabulary size to the size recommended by VOLT (Xu et al., 2020), then lowering the `max_sentencepiece_length` until the average downsampling factor is close to 4.⁸ We chose 4 as it is roughly equivalent to the ratio of the number of bytes per subword token when comparing the ByT5 and mT5 models across all the languages in the mC4 corpus (Xue et al., 2022).

Our parameter setup for models, vocabulary, and training are noted in Appendix B.

6 Results

We start by comparing the 4 downsampling methods, discerning the importance of the positional, length, and morpheme consistency. We then compare our best-performing downsampling method to our two standard baselines: the subword-level and character-level models. Third, we compare these models in their ability to generalize to out-of-domain datasets. Lastly, we look at the encoder-only tasks, comparing all methods thus far.

6.1 Comparing Downsampling Methods

First comparing the various downsampling methods Table 3, we see that as expected from Section 3 our novel SDD method performs best overall.

Of the three factors, positional consistency appears most important, making up the greatest portion of the performance increase. Given that

⁷By default, it is set to 16, which inhibits the performance of the LSTM in the decoder.

⁸The average downsampling factor is calculated by averaging the number of bytes per subword token in the training data.

Method	BLEU							COMET						
	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg
Subword	27.23	24.08	25.59	11.22	14.73	11.30	19.02	0.203	0.073	0.063	0.120	-0.097	0.176	0.090
Char	27.37	24.32	26.34	8.73 [†]	15.52 [*]	11.87[*]	19.03	0.222	0.040	0.142 [*]	-0.008 [†]	-0.026 [*]	0.190	0.093
SDD	27.78	25.12[*]	27.74[*]	11.95[*]	15.99[*]	11.31	19.98	0.237[*]	0.023[*]	0.213[*]	0.146	-0.006[*]	0.175	0.131

Table 4: Translation results of traditional subword models and character models without downsampling compared to character models with subword-delimited downsampling (SDD). Green (*) and red (†) denote a significant positive or negative difference ($p < 0.05$) with respect to the Subword model.

positional inconsistency means there is often no subword-like structure to the downsampled tokens, it is understandable that the model has the most difficulty translating in this scenario.

Length consistency contributes some improvement as well, showing that there is indeed a bottleneck effect when downsampling longer words, and that there is merit in the model splitting such words into multiple tokens prior to the Transformer.

Morpheme consistency is least important; however it seems to have a larger impact on the more morphological languages of Arabic and Turkish. This gives evidence towards the idea that the subword splitting of SentencePiece is likely more conducive to translating morphemes than translating 4-character chunks. Although we use SentencePiece due to its ubiquity, there are arguably better subword tokenizers with respect to adherence to morphology, such as LMVR (Ataman et al., 2017). These tokenizers are fully compatible with SDD, and they may further increase the disparity we see in morpheme consistency.

6.2 Comparing to Subword and Character Models

We now compare the performance of SDD to subword and character models in Table 4.

Overall, SDD outperforms both subword and character models on both BLEU and COMET. It performs significantly better than the subword model in 4 and 3 cases according to BLEU and COMET, respectively. Only in 1 case in terms of COMET does SDD lead to significantly worse performance.

The Arabic–English language pair shows the largest improvement over the baselines. The reason for this is unclear, though it is the only language pair with an entirely different character set between the source and target (save for numbers, symbols, and some proper nouns).

As SDD was originally intended to shore up the weaknesses of the previous downsampling meth-

ods, the aim was to perform on par with the subword and character models. The slight increase in performance shows that there is some benefit in using a combination of the two, namely using character-level input while operating on a subword-level within the Transformer encoder-decoder.

6.3 Out-of-domain Generalization

Boukkouri et al. (2020) found that with CharacterBERT, the models generalized better to out-of-domain encoder-only tasks such as classification of medical data. They argue that because CharacterBERT does not have a strict vocabulary, and it learns more general properties of language which can be useful for unseen words. As we see evidence of the same in our embedding analysis (see Section 7), we similarly test our models on two out-of-domain translation datasets. For all languages, we test on FLoRes (Goyal et al., 2022), which consists of Wikipedia data. For English–German, we additionally test on the WMT21 Biomedical Shared Task test set.

The FLoRes and biomedical results are shown in Tables 5 and 6, respectively. Here we see the SDD model and character model both outperforming the subword model, as expected. The results on FLoRes favor the SDD model, while the biomedical results favor the character model. The biomedical data is arguably “more out-of-domain” than FLoRes, since it contains a large amount of medical terminology unlikely to appear in the training data. This may indicate that the character downsampling is somewhat sensitive to the vocabulary it is trained on, and as such it generalises better than the subword model but not as well as the character model.

Since the SDD model only uses a subword vocabulary to determine the lengths for downsampling, it is perhaps possible to use a different vocabulary when switching domains, namely one that offers better downsampling for the out-of-domain words. We leave this for future research.

Method	BLEU							COMET						
	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg	de-en	en-de	ar-en	en-ar	tr-en	en-tr	avg
Subword	19.88	16.79	13.90	9.07	16.25	11.02	14.49	-0.158	-0.385	-0.344	-0.292	-0.053	-0.031	-0.210
Char	21.97	17.46	15.86	8.95	17.59	12.03	15.64	0.003	-0.297	-0.132	-0.340	0.065	0.049	-0.109
SDD	21.45	17.46	16.75	10.73	17.85	12.01	16.04	-0.065	-0.330	-0.094	-0.154	0.074	0.035	-0.089

Table 5: Results on FLoRes evaluation set.

Method	BLEU		COMET	
	de-en	en-de	de-en	en-de
Subword	9.37	4.36	-0.729	-1.119
Char	11.71	4.74	-0.499	-0.947
SDD	9.41	4.05	-0.659	-1.036

Table 6: Results on Biomedical evaluation set.

6.4 Encoder-only Tasks

The accuracies achieved in the encoder-only tasks are shown in Table 7. The character-level model has surprisingly low accuracies. Given that the main difference between the character model and the other models is the longer sequence length fed into the Transformer, we expect the complexity of the self-attention patterns necessary for these tasks is more difficult to learn when data is limited.

Method	NLI (%)	RC (%)
Subword	80.14	73.90
Char	60.75	68.86
Fixed	81.01	73.28
Buf. Fixed	81.81	73.05
WDD	81.89	73.69
SDD	81.30	74.22

Table 7: Results of NLI and review classification (RC).

Our SDD method outperforms both subword and character baselines on both tasks. Unlike with translation, the downsampling models show little difference in performance. As SDD is intended to help by providing a consistent tokenization, it seems that this consistency is less important for sequence classification tasks. This is probably because there is no need for character or word recovery: the model does not need to reconstruct any of its input, so it can potentially lose some character information while still learning to correctly classify.

A token classification task such as part-of-speech tagging may be more difficult for the fixed-size downsampling model, although it is not clear how to apply such a model to a token classification task, given the mismatch in downsampling blocks and token labels. The SDD model can however be applied to token classification in the same manner a subword model would be applied.

Another reason for there being more of a gap in the performance for translation might be that the cross-attention is what benefits most from the three token consistency factors. Inconsistent tokens on both the source and target side likely make learning what to attend to quite difficult.

7 Embedding Analysis

The sole difference in the architecture of the encoder-only models is their embeddings, or the hidden representations prior to being fed into the Transformer. As such, we extract these embeddings for words to analyze their differences. We compare the models in their word embedding similarity (i.e. cosine similarity) for pairs of words. We generate 4 test sets:

1. Grammatical pairs - Pairs which share a common lemma.
2. Close pairs - Pairs with a Levenshtein distance of 1.
3. Far pairs - Pairs with a max Levenshtein distance.
4. Far Synonyms - Synonyms with a max Levenshtein distance.

The first two sets are meant to distinguish beneficial character-level similarity (e.g. “take” and “takes”) from detrimental similarity (e.g. “pour” and “tour”). The latter two sets are meant to distinguish the opposite, when a lack of similarity may be appropriate (e.g. “bulb” and “faster”) or inappropriate (e.g. “camp” and “tent”). We limit the vocabulary used in the test sets to words in the subword vocabulary that only require a single token. We also filter out tokens which have no synonyms, effectively also removing non-words and non-English words. Lemmas and synonyms are gathered using WordNet (Miller, 1995).

We split the test sets additionally into 3 sub-categories: “seen”, where both words in the pair have been seen during training, “half-seen”, where only one word is seen in training, and “unseen”,

		Grammatical	Close Spell	Far Spell	Far Synonym
Subword	Seen	1.03 ± 0.08	0.17 ± 0.06	0.05 ± 0.05	0.18 ± 0.05
μ : 0.01	Half-seen	-0.09 ± 0.07	-0.09 ± 0.02	-0.12 ± 0.02	-0.12 ± 0.02
σ : 0.08	Unseen	-0.12 ± 0.12	-0.10 ± 0.05	-0.10 ± 0.02	-0.06 ± 0.09
SDD	Seen	5.75 ± 0.05	4.68 ± 0.06	-0.52 ± 0.04	-0.49 ± 0.04
μ : 0.31	Half-seen	6.09 ± 0.09	5.17 ± 0.06	-0.43 ± 0.04	-0.40 ± 0.04
σ : 0.09	Unseen	6.06 ± 0.14	5.62 ± 0.11	-0.27 ± 0.04	-0.21 ± 0.13

Table 8: Model average z -scores for the respective test sets, with a 95% confidence interval.

where neither is seen in training.⁹ The sizes of each split are shown in Table 12 in Appendix A. The results are shown in Table 8.

As expected, the subword model has near 0 similarity for half-seen and unseen words since it has no mechanism for developing their embeddings during training. It does recognize grammatical seen words as similar, but it has a more muted response to seen words with similar spelling. In contrast, SDD has a strong response to both seen and unseen words, as the character models can develop embeddings for unseen words according to character-level patterns. We see that both models appear capable of distinguishing between grammatical pairs and pairs with similar spelling but not necessarily similar meaning, though only in the case where both words are seen for the subword model.

There is a stark difference between the z -scores of the Close Spell and Far Spell sets for the SDD model. As for the Far Synonyms, there is a small difference for the subword model, and none for the SDD model. It is possible that these words are distinguished in later layers of the model.

8 Practicality Discussion

While this paper takes a more theoretical approach, using smaller, domain-specific datasets for training NMT, the practical usage is worth considering. Given the prior work of ByT5, Charformer, and CharacterBERT, the consensus appears that with large, pretrained models using character information, the performance on standard metrics is similar, and the outperformance is on out-of-domain data (Boukkouri et al., 2020) or data with character-level corrupted input (Xue et al., 2022). As such, we expect the same is true for SDD, however the upsampling module used (based on Libovický et al. (2021)) appears to be a limiting factor. We explain

⁹For these experiments, we used separate models trained using the same initial vocabulary as T5-Small, in order to get a larger number of single-token words to evaluate on. This is why there exist words in the vocabulary that are not seen during training. The performance of these models is similar to those reported in Table 7.

this in further detail in Appendix C, where we train and test the models on the larger WMT14 German-English dataset, finding that the performance of the subword model deteriorates when the same two-step decoding method is added.

Conversely, we expect such models to be useful in lower-resource settings. In such an scenarios, we show that the inclusion of character-level information improves performance beyond that of subword models (Appendix D).

	Translation			Encoder Tasks		
	Iter	Epochs	Total	Iter	Epochs	Total
Char	×3.58	×1.31	×4.69	×2.23	×2.04	×4.55
SDD	×2.61	×1.28	×3.34	×1.34	×1.36	×1.82

Table 9: Training time ratios with respect to the subword model. “Iter” refers to the iteration speed (e.g. Char is 3.58 times slower than the subword model per iteration for translation), “Epochs” refers to the total number of epochs needed, and “Total” is the product of the first two.

In terms of training time, we report the training time ratios with respect to the subword model in Table 9.¹⁰ While the SDD model is considerably slower than the subword model in terms of training time, it still performs better than the character model, particularly for encoder-only tasks, suggesting that decoding is the main source of the slower training times.

One potential use of the SDD model that is yet unexplored comes in the form of adapting existing subword models to take character input. Since pre-training is a costly endeavor, and since there are beneficial characteristics of character-level models such as out-of-domain generalization, it may be possible to improve a subword model by adapting it to use the SDD downsampling module rather than its own subword embeddings. Previous work in adaptation has shown success in adapting models to different tasks (Üstün et al., 2020; Pfeiffer et al., 2020) and languages (Bapna et al., 2019; Artetxe

¹⁰All experiments were conducted on a single Nvidia V100 GPU.

et al., 2019; Üstün et al., 2021), so a similar approach may be useful here.

9 Conclusion

Previous work has casted doubt on the usefulness of character-level NMT models, due to their lack of improvement over subword models despite using more fine-grained information while being also slower. Downsampling modules added to the character models have previously been proposed but have always come at the cost of accuracy.

We show that it is possible to downsample without sacrificing accuracy, by downsampling based on the lengths of subwords. This novel downsampling outperforms the previous downsampling methods, as well as character and subword models on the majority of language pairs tested.

There are several avenues for future research. While much work has been done on optimizing a vocabulary for a subword model, finding the optimal lengths for subword-delimited downsampling is still an open problem. The most promising may in fact be adapting any of the numerous pretrained subword models to use characters as input. Overall, character-level models show promise that has yet to be fully realized.

10 Limitations

As we mention in Sections 5 and 8, our largest limitation is that the majority of our testing is done on smaller datasets, consisting of roughly 200 thousand sentence pairs per language pair. While we do test on a larger dataset for German→English, it is limited to only a single language pair, and single direction, so it is still an open question whether our method works well in general in higher-resource settings.

Additionally, we only test on 3 language pairs in our main results, all of which have English on the source or target side. It is possible that our method only works well in circumstances where English is present in the language pair, or it only works well where the other language is either German, Arabic, or Turkish. Notably, languages such as Chinese can have characters that hold the same meaning as a word in English, and as such subword tokenizers like SentencePiece may be less useful. Subsequently, SDD may be less useful for these types of languages.

In terms of the tasks tested, we mainly focus on translation, so we can make no claims about the

performance on other generative tasks. We test on two discriminative tasks, NLI, and review classification, however both are tested without following the popular pretrain-then-finetune paradigm, making the results difficult to compare to existing work. The scope of these tests is also limited to English only.

In terms of parameters tested, we mainly follow previous work, so it is possible that our method does not perform as well (or possibly performs better) under different parameter settings. To keep our carbon impact minimal, we opted for using VOLT to determine an optimal vocabulary size (which requires no training of NMT models), rather than the standard grid search approach. VOLT does not guarantee an optimal vocabulary size however, and this may have an impact on our results, be it favorably or unfavorably.

Finally, as we note in Appendix B.3, our method is considerably slower than the subword model. However, it is also considerably faster than the character model while also performing better on the majority of the tasks tested. As this work is exploring new territory, it is very likely that our implementation is not as efficient as it could be.

11 Acknowledgements

We thank the Center for Information Technology of the University of Groningen for their support and for providing access to the Peregrine high performance computing cluster.

References

- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2019. On the cross-lingual transferability of monolingual representations. *arXiv preprint arXiv:1910.11856*.
- Duygu Ataman, Matteo Negri, Marco Turchi, and Marcello Federico. 2017. Linguistically motivated vocabulary reduction for neural machine translation from turkish to english.
- Ankur Bapna, Naveen Arivazhagan, and Orhan Firat. 2019. Simple, scalable adaptation for neural machine translation. *arXiv preprint arXiv:1909.08478*.
- Hicham El Boukkouri, Olivier Ferret, Thomas Lavergne, Hiroshi Noji, Pierre Zweigenbaum, and Junichi Tsujii. 2020. Characterbert: Reconciling elmo and bert for word-level open-vocabulary representations from characters. *arXiv preprint arXiv:2010.10392*.
- Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated

- corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.
- Marta R Costa-Jussa and José AR Fonollosa. 2016. Character-based neural machine translation. *arXiv preprint arXiv:1603.00810*.
- Lukas Edman, Antonio Toral, and Gertjan van Noord. 2022. Patching leaks in the charformer for efficient character-level generation. *arXiv preprint arXiv:2205.14086*.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’ Aurelio Ranzato, Francisco Guzman, and Angela Fan. 2022. The flores-101 evaluation benchmark for low-resource and multilingual machine translation. *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Jason Lee, Kyunghyun Cho, and Thomas Hofmann. 2017. Fully character-level neural machine translation without explicit segmentation. *Transactions of the Association for Computational Linguistics*, 5:365–378.
- Jindřich Libovický, Helmut Schmid, and Alexander Fraser. 2021. Why don’t people use character-level machine translation? *arXiv preprint arXiv:2110.08191*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Myle Ott, Sergey Edunov, David Grangier, and Michael Auli. 2018. *Scaling neural machine translation*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. *Bleu: a method for automatic evaluation of machine translation*. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. Mad-x: An adapter-based framework for multi-task cross-lingual transfer. *arXiv preprint arXiv:2005.00052*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. *COMET: A neural framework for MT evaluation*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Yi Tay, Vinh Q Tran, Sebastian Ruder, Jai Gupta, Hyung Won Chung, Dara Bahri, Zhen Qin, Simon Baumgartner, Cong Yu, and Donald Metzler. 2021. Charformer: Fast character transformers via gradient-based subword tokenization. *arXiv preprint arXiv:2106.12672*.
- Ahmet Üstün, Alexandre Berard, Laurent Besacier, and Matthias Gallé. 2021. Multilingual unsupervised neural machine translation with denoising adapters. *arXiv preprint arXiv:2110.10472*.
- Ahmet Üstün, Arianna Bisazza, Gosse Bouma, and Gertjan van Noord. 2020. Uadapter: Language adaptation for truly universal dependency parsing. *arXiv preprint arXiv:2004.14327*.
- Daimeng Wei, Zongyao Li, Zhanglin Wu, Zhengzhe Yu, Xiaoyu Chen, Hengchao Shang, Jiaxin Guo, Minghan Wang, Lizhi Lei, Min Zhang, et al. 2021. Hwts’s participation in the wmt 2021 news translation shared task. In *Proceedings of the Sixth Conference on Machine Translation*, pages 225–231.
- Jingjing Xu, Hao Zhou, Chun Gan, Zaixiang Zheng, and Lei Li. 2020. Vocabulary learning via optimal transport for neural machine translation. *arXiv preprint arXiv:2012.15671*.
- Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. Byt5: Towards a token-free future with pre-trained byte-to-byte models. *Transactions of the Association for Computational Linguistics*, 10:291–306.

A Datasets

The datasets used for translation are shown in Table 10. In summary, we use IWSLT2017,¹¹ SETIMES,¹² and datasets provided by WMT shared tasks.¹³

The datasets used for NLI and review classification (RC) are shown in Table 11. For NLI, we use the SNLI corpus (Bowman et al., 2015), and for RC we use the Amazon Customer Reviews Dataset.¹⁴ For SNLI we use the predefined train/dev/test splits, and for RC we take a 99/1 split of the Books_v1_00 subset for training and validation, and use 1% of the Books_v1_02 subset for testing (amounting to roughly 22k test sentences).

Concerning the evaluation sets we generated for Section 7, we show the sizes of the generated sets in Table 12. We limited the size of the test sets to 2000 word pairs, randomly selecting from those generated.

¹¹<https://sites.google.com/site/iwslt2017/TED-tasks>

¹²<https://opus.nlpl.eu/SETIMES.php>

¹³<https://statmt.org/wmt21/translation-task.html>

¹⁴<https://s3.amazonaws.com/amazon-reviews-pds/readme.html>

Lang pair	Main			Low-Resource	High-Resource
	en \leftrightarrow de	en \leftrightarrow ar	en \leftrightarrow tr	xh \leftrightarrow zu	de \rightarrow en
Train	IWSLT2017	IWSLT2017	SETIMES	CC-Aligned	WMT14
Size (k)	199	224	197	85	4509
Eval	IWSLT2017	IWSLT2017	WMT17/18 dev	WMT21 dev	WMT13/14 dev
Domain	TED Talks	TED Talks	News	Crawl	News / Crawl / Govt

Table 10: Datasets used for translation, showing the names of the training and evaluation sets, the size of the training data (in sentence pairs), and the domain.

Task	NLI	RC
Train	SNLI train	Books_v1_00
Size (k)	547	5711
Eval	SNLI dev/test	Books_v1_00 / 02
Domain	Image captions	Book reviews

Table 11: Datasets used for encoder-only tasks, showing the names of the training and evaluation sets, the size of the training data (in premise/hypothesis pairs or reviews), and the domain.

	Grammatical	Close spell	Far spell	Far Synonym
Seen	1873	2000	2000	2000
Half-seen	273	2000	2000	2000
Unseen	92	525	2000	164

Table 12: Size of generated datasets for evaluating the embeddings of word pairs.

B Parameters

B.1 Architecture

For the CNN downsampling module, we use the same parameters as Libovický et al. (2021).

We use Transformer Base for our Encoder-Decoder, only modifying the maximum position embeddings to 2048 to account for the increased length of our character models.

For the upsampling module, we use the same parameters as used in Libovický et al. (2021), with the exception of the linear layer which projects the hidden representations of the Transformer decoder to the LSTM. The output size used in Libovický et al. (2021) is 64 times the downsampling factor. For variable-length downsampling, this needs to be set to the maximum word length in bytes plus one (accounting for the end-of-word token).

	de \leftrightarrow en	ar \leftrightarrow en	tr \leftrightarrow en	xh \leftrightarrow zu
Subword	62M	55M	51M	50M
Char	46M	46M	46M	46M
SDD	56M	56M	56M	56M

Table 13: Number of parameters (in millions) for each model trained.

The total number of parameters in each model is

shown in Table 13. We opted to keep the number of parameters in the Transformer fixed, meaning the variation in parameter count comes only from the embeddings as well as the downsampling and upsampling modules, if they are present. The number of parameters varies across language for the subword model due to the varying size of the vocabulary.

B.2 Vocabulary

Lang pair	Main			LR	HR
	en \leftrightarrow de	en \leftrightarrow ar	en \leftrightarrow tr	xh \leftrightarrow zu	de \rightarrow en
Vocab size (k)	16	9	5	4	16
Max token len	6	4	6	5	6

Table 14: Vocabulary sizes and max token lengths for the languages tested.

As noted in Section 5, we use VOLT to determine the optimal vocabulary sizes. We specify these, along with the max token length that most closely matches an average downsampling factor of 4 in Table 14. It should be noted that the max token length argument in SentencePiece is determined by the maximum character length, however SDD operates on the byte level. So for English–Arabic, although the max token length is set to 4, the average downsampling factor is in fact above 4, due to the length of the Arabic characters in UTF-8 being typically 2-3 bytes.

B.3 Training

We use an effective batch size of 128, accumulated over 4 iterations. We use the AdamW optimizer (Loshchilov and Hutter, 2017), with betas 0.9 and 0.999, a learning rate of $2e-4$, and a linear scheduler with a warmup of 10000 steps. We apply an early stopping with a patience of 10, using the BLEU score on the validation set as the stopping criterion.

C Higher-Resource Translation

We also run experiments in the higher resource setting of the WMT14 German \rightarrow English data. Fol-

lowing the principles of Ott et al. (2018), we use larger batch sizes of 50k and 240k tokens for the subword and character models, respectively, both of which average to about 2000 sentences per batch. We also increase the learning rate to 5e-4 and use mixed-precision for the training. We include an additional model, which is the subword model with the two-step decoding as used in the models with downsampling. In other words, the subword model is given the same LSTM upsampling head, but the upsampling is simply 1-to-1. We test this to ablate any effect of the upsampler on the results.¹⁵

	BLEU	COMET
Subword	28.62	0.431
Two-step subword	27.49	0.360
SDD	27.27	0.364

Table 15: Translation results on the WMT14 DE→EN dataset.

Our results are shown in Table 15. We can see that the standard subword model performs best, however the addition of the two-step decoder hurts performance significantly. The subword model with the two-step decoder still performs on par with our SDD model, confirming our expectations. We conclude that SDD is competitive in higher-resource settings, since it achieves similar scores to the subword variant that also has a two-step decoder. This two-step decoder, while effective on lower-resource settings, does not scale well to higher-resource settings without any modification. This raises the question as to whether an upsampling method exists that has better scaling ability. We leave this for future research to explore.

D Lower-Resource Translation

To analyze our models on lower-resource translation, we chose to train and evaluate on Xhosa–Zulu, specifically the data provided by the WMT2021 Shared Task. Typically, in lower-resource scenarios, other techniques such as multilingual transfer learning and back-translation are applied to improve the models. We chose this language pair as the performance gain from employing these techniques is less substantial (Wei et al., 2021).

The results show that the character-level model without downsampling performs best, with SDD as a close second. Xhosa and Zulu, being closely related languages, likely would benefit greatly from

¹⁵We do not include this in our main results because we did not see a noticeable difference in the performances there.

Method	BLEU		COMET	
	xh-zu	zu-xh	xh-zu	zu-xh
Subword	4.59	5.89	-0.3222	-0.1149
Char	6.51	6.42	-0.0071	0.0286
SDD	6.35	6.32	-0.1282	-0.0285

Table 16: Translation Results for Xhosa–Zulu.

character-level translation regardless of the amount of data, given that many of the differences between the two are on the character level. Of course, more thoroughly evaluating SDD on low-resource settings would require a pretrain-then-finetune approach, which we reserve for future research.