

Joint Optimization of Tokenization and Downstream Model

Tatsuya Hiraoka[†], Sho Takase[†], Kei Uchiumi[‡], Atsushi Keyaki[‡], Naoaki Okazaki[†]

[†] Tokyo Institute of Technology

[‡] Denso IT Laboratory, Inc.

{tatsuya.hiraoka, sho.takase}@nlp.c.titech.ac.jp

{kuchiumi, akeyaki}@d-itlab.co.jp

okazaki@c.titech.ac.jp

Abstract

Since traditional tokenizers are isolated from a downstream task and model, they cannot output an appropriate tokenization depending on the task and model, although recent studies imply that the appropriate tokenization improves the performance. In this paper, we propose a novel method to find an appropriate tokenization to a given downstream model by jointly optimizing a tokenizer and the model. The proposed method has no restriction except for using loss values computed by the downstream model to train the tokenizer, and thus, we can apply the proposed method to any NLP task. Moreover, the proposed method can be used to explore the appropriate tokenization for an already trained model as post-processing. Therefore, the proposed method is applicable to various situations. We evaluated whether our method contributes to improving performance on text classification in three languages and machine translation in eight language pairs. Experimental results show that our proposed method improves the performance by determining appropriate tokenizations.

1 Introduction

Tokenization, which converts a raw sentence into a sequence of tokens, is a crucial process that affects the performance of NLP tasks. Existing studies have proposed various tokenization methods including rule-based tokenization (Koehn et al., 2007), dictionary-based tokenization (Kudo, 2006; Morita et al., 2015; Tolmachev et al., 2018; Takaoka et al., 2018), supervised tokenization with neural networks (Yang et al., 2017; Cai et al., 2017; Yang et al., 2019), and unsupervised tokenization (Goldwater et al., 2006, 2009; Mochihashi et al., 2009; Sennrich et al., 2016; Kudo and Richardson, 2018). Much of prior research has reported that an appropriate tokenization depends on each downstream task (Xu et al., 2008; Chang et al., 2008; Nguyen

et al., 2010; Domingo et al., 2018; Hiraoka et al., 2019; Gowda and May, 2020). Moreover, Hiraoka et al. (2020) implies that we have to consider a downstream model to determine an appropriate tokenization. In other words, we can improve the performance of a downstream model by determining an appropriate tokenization for the downstream model. However, since traditional tokenizers are isolated from a downstream model, we need to train a given downstream model with each possible tokenization and evaluate its performance to determine the appropriate tokenization. Performing such an exploration whenever we construct a new downstream model is impractical.

Several studies have addressed the optimization of a tokenizer based on a downstream task or/and model (He et al., 2020; Hiraoka et al., 2020), but existing methods are restricted to specific tasks. He et al. (2020) proposed DPE as a tokenization method for a sequence-to-sequence problem such as machine translation. Their method trains a tokenizer with a given training corpus, but it is isolated from a downstream model such as a neural encoder–decoder for machine translation. Hiraoka et al. (2020) proposed OpTok, which jointly trains a tokenizer and a downstream model. However, its architecture is specific to classification problems based on sentence representations, and thus, it cannot be applied for various tasks such as sequence-to-sequence problems. Therefore, there is no method to optimize a tokenizer depending on any downstream task and model.

In this paper, we propose a novel method to jointly optimize a tokenizer and downstream model without any restriction on a task¹. The proposed method can determine an appropriate tokenization for a downstream model because it explores different tokenizations based on loss values of the down-

¹Code: <https://github.com/tatHi/optok4at>

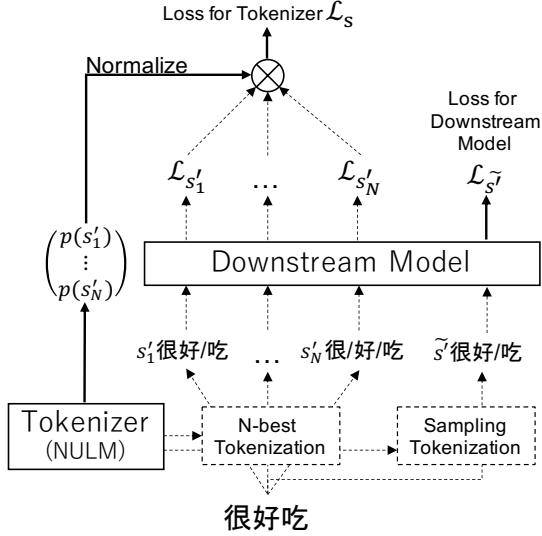


Figure 1: Overview of the proposed method calculating a loss for a tokenizer \mathcal{L}_s and that for a downstream model $\mathcal{L}_{\tilde{s}}$ using N -best tokenizations (Section 2.2) and a sampled tokenization (Section 2.3), respectively. The arrows with the continuous lines indicate the differentiable path for back-propagation.

stream model. Since the proposed method requires only loss values of the downstream model, we can apply it for any task and model. Moreover, even if a given downstream model is already trained, our proposed method can be applied to improve the performance by refining tokenization. We call this refinement of tokenization *post-processing*. Thus, we can easily use the proposed method in various situations including the case where we have a sufficiently trained downstream model.

We conducted experiments on text classification and machine translation tasks in various languages. Experimental results indicate that the proposed method outperformed existing tokenization methods in both the tasks. We also showed that our method can enhance the performance by refining tokenization as post-processing for downstream models trained with subword regularization (Kudo, 2018; Provilkov et al., 2020).

2 Proposed Method

The proposed method comprises a tokenizer and a downstream model. We optimize the two modules simultaneously. First, we present the training outline of the case where we use one sentence as an input in Section 2.1. Second, we introduce the training of the tokenizer (Section 2.2) and the downstream model (Section 2.3). Finally, we explain the training strategy for a task that requires multiple inputs such as machine translation (Section 2.4).

2.1 Optimizing Tokenization with Loss

The proposed method tokenizes a sentence s into a sequence of words w in vocabulary V , $s' = w_1, \dots, w_I$, where I is the sequence length. In this tokenization process, our purpose is to minimize the following loss value:

$$\mathcal{L}_{s'} = q(f(s'), z), \quad (1)$$

where $f(s')$ is a downstream model that outputs a prediction of the downstream task from a tokenized sentence s' , and $q(f(\cdot), z)$ is a task-specific loss function between a model prediction and supervisory signal z .

Figure 1 presents an outline of the proposed method. To determine the tokenization satisfying $\text{argmin}_{s'}(q(f(s'), z))$, we update the tokenizer to assign a higher probability to a useful tokenization for the downstream model. Concretely, we construct N tokenizations $s'_1, \dots, s'_n, \dots, s'_N$ for a training instance and then compute loss values for each tokenization. We weight each loss based on probability $p(s'_n)$ computed by the tokenizer and use the weighted sum to train the tokenizer, as follows:

$$a_n = \frac{p(s'_n)}{\sum_{m=1}^N p(s'_m)}, \quad (2)$$

$$\mathcal{L}_s = \sum_{n=1}^N a_n \mathcal{L}_{s'_n}. \quad (3)$$

In this study, we used N -best tokenizations. In these equations, we weight losses $\mathcal{L}_{s'_1}, \dots, \mathcal{L}_{s'_N}$ corresponding to N -best tokenizations with their sentence probabilities normalized such that the sum is 1. By optimizing the tokenizer based on the weighted sum, \mathcal{L}_s , the tokenizer assigns high probability to the appropriate tokenization for the downstream model.

We can use any function for $f(\cdot)$ and $q(f(\cdot), \cdot)$ in Eq.(1). Therefore, the proposed method has no restrictions on the downstream task and model. For instance, in the case where text classification is the downstream task, $f(\cdot)$ is a neural network predicting a label of a given tokenized sentence, and $q(f(\cdot), \cdot)$ is the cross-entropy loss between the model prediction and the true label.

2.2 Tokenizer: NULM

We employ a neural unigram language model (NULM) as our tokenizer. It calculates the unigram probability of a word $p(w)$ with a word embedding

v_w as follows:

$$d_w = \text{MLP}(v_w), \quad (4)$$

$$p(w) = \frac{\exp(d_w)}{\sum_{\tilde{w} \in V} \exp(d_{\tilde{w}})}, \quad (5)$$

where $\text{MLP}(\cdot)$ is a multilayer perceptron. We initialize vocabulary V with a reasonable size of words. For example, we can use a tokenization by SentencePiece (Kudo and Richardson, 2018) or BPE (Sennrich et al., 2016) for the initialization. We also calculate the probability of a tokenization $p(s')$ as follows:

$$p(s') = \prod_{w \in s'} p(w). \quad (6)$$

For the training with Eq.(3), we obtain N -best tokenizations by applying Forward-DP Backward-A* algorithm (Nagata, 1994) for possible tokens against sentence s . In the inference phase, we can also obtain the 1-best appropriate tokenization for the downstream task using Viterbi algorithm (Viterbi, 1967) for the trained NULM.

2.3 Downstream Model Training

We can train the downstream model with loss \mathcal{L}_s in Eq.(3), but we use subword regularization (Kudo, 2018) to obtain a better model. Thus, we compute $\mathcal{L}_{\tilde{s}'} = q(f(\tilde{s}'), z)$ for a sampled tokenization \tilde{s}' and use $\mathcal{L}_{\tilde{s}'}$ to train the downstream model.

We sample tokenization \tilde{s}' from $p(\tilde{s}')^\alpha / \sum_{k=1}^K p(s'_k)^\alpha$ computed by the NULM in Eq.(5) (Kudo, 2018). Here, $\alpha \in \mathbb{R}^+$ is a hyperparameter that controls the diversity of the sampled tokenization. If we set α as a lower value, the distribution is similar to the uniform distribution; otherwise, the distribution strongly depends on each tokenization probability $p(s')$. K is also a hyperparameter denoting the number of candidates for sampling, and we use Forward Filtering Backward Sampling (Scott, 2002; Mochihashi et al., 2009) if $K = \infty$.

Subword regularization not only sophisticates the downstream model but also provides various tokenizations to the downstream model during training. Therefore, subword regularization helps in exploring the appropriate tokenization.

2.4 Training in Multiple Sentences as Inputs

Previous sections discussed the case where we use one sentence as an input, but we have to input multiple sentences to the downstream model in

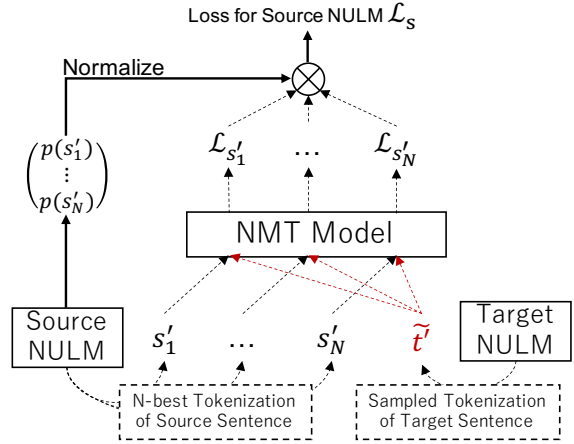


Figure 2: Overview of the calculation of a tokenization loss \mathcal{L}_s for the source-side NULM in NMT requiring two inputs, source and target sentences s and t , respectively. The arrows with the continuous line indicate the differentiable path for back-propagation.

some tasks. This section describes our training strategy in such cases.

To compute the loss value for training the tokenizer, we consider multiple tokenizations for one sentence and use the sampled tokenization for the others. For example, in machine translation, we input the source and target sentences to train the downstream model. The source sentence is the input of the downstream model, and the target sentence is the supervisory signal. Let s and t be the source sentence and target sentences, respectively, and s' and t' be the corresponding tokenizations. We update the NULM of the source side using $\mathcal{L}_{s'_n} = q(f(s'_n), t')$, where t' is a sampled tokenization for the target sentence. We also compute the loss for the NULM of the target side with $\mathcal{L}_{t'_n} = q(f(\tilde{s}'), t'_n)$, where \tilde{s}' is a sampled tokenization. For training the downstream model, we use sampled tokenizations for all the input sentences. Thus, we compute $\mathcal{L}_{\tilde{s}', \tilde{t}'} = q(f(\tilde{s}'), \tilde{t}')$ for the downstream model. We outline this training process for the NULM of the source side in Figure 2, and the training for the target side is explained in the same manner.

3 Experiment

To validate the applicability of the proposed method to various downstream tasks, we conducted experiments on text classification and machine translation tasks from existing literature. To compare our method with the existing methods that determine the appropriate tokenization for a specific downstream task, we employ OpTok (Hiraoka

	SP	SP+R	OpTok	Ours
Weibo(Zh)*	92.70	92.79	92.82	93.06
Twitter(Ja)*	85.89	86.51	86.97	86.92
Twitter(En)*	75.98	77.31	78.52	78.88
Genre(Zh)	44.19	47.95	48.18	48.41
Rating(Zh)	48.96	49.41	49.63	49.76
Genre(Ja)	46.82	49.84	50.15	50.79
Rating(Ja)	52.88	53.43	53.55	53.69
Genre(En)*	69.93	71.68	71.88	71.83
Rating(En)*	66.42	67.53	67.68	67.90
SNLI*	75.62	76.75	77.04	77.05

Table 1: Experimental results on text classification tasks (F1-score). SP and R indicate SentencePiece and subword regularization, respectively. The highest scores are highlighted in bold. Scores of SP+R and OpTok on datasets with * are quoted from Hiraoka et al. (2020).

et al., 2020) for text classification and DPE (He et al., 2020) for machine translation.

3.1 Text Classification

Settings We utilized ten datasets of text classification tasks in three languages. Weibo(Zh)², Twitter(Ja)³, and Twitter(En)⁴ are sentiment analyses on SNS corpora in Chinese, Japanese, and English, respectively. Genre and Rating are genre prediction and rating predictions from reviews posted on E-commerce corpora, respectively, in Chinese (Zhang et al., 2015)⁵, Japanese (Rakuten, Inc., 2014), and English (He and McAuley, 2016)⁶. In addition, we employed the SNLI corpus (Bowman et al., 2015) to evaluate our method on the setting requiring two sentences as the input.

We focus on SentencePiece (SP) (Kudo and Richardson, 2018) and OpTok (Hiraoka et al., 2020) as other tokenizers for comparison with the proposed method. OpTok is a method to optimize tokenization for text classification by weighting a sentence vector with N -best tokenization. In addition, we trained each model with subword regularization (SP+R) (Kudo, 2018) for fair comparisons. In subword regularization, we used a sampled tokenization for the training phase and a 1-best tokenization for the inference phase.

For the downstream model, we used a BiLSTM encoder since we followed the experimental configurations of OpTok⁷. We used SentencePiece to

²<https://github.com/wansho/senti-weibo>

³http://www.db.info.gifu-u.ac.jp/data/Data_5d832973308d57446583ed9f

⁴<https://www.kaggle.com/c/twitter-sentiment-analysis2>

⁵<http://yongfeng.me/dataset/>

⁶<http://jmcauley.ucsd.edu/data/amazon/>

⁷Our implementation is based on the existing code: <https://github.com/tatHi/optok>

construct vocabulary V^8 and initialized unigram probabilities of our NULM. The initial vocabulary sizes are 16K for Twitter(Ja) and Twitter(En) and 32K for the others. The number of tokenizations is $N = 3$, and the hyperparameters for subword regularization are $\alpha = 0.2$ and $k = \infty$. For the SNLI corpus, the system shares the same NULM for the premise and hypothesis and optimizes the NULM in the manner explained in Section 2.4.

Results Table 1 presents the experimental results on text classification. This table indicates that our proposed method surpasses OpTok in eight datasets. For the other two datasets, the performance of our method is comparable to OpTok. These results indicate that the proposed method is a better alternative to the existing tokenizers for text classification tasks. We consider that the difference in the performance is caused by the difference in strategy between the training the downstream models. OpTok trains the downstream model with a weighted sum of sentence vectors corresponding to N -best tokenization with their tokenization probabilities, but it uses 1-best tokenization in the inference. This gap might harm the downstream model. In contrast, since our method trains the downstream model with only one sampled tokenization, the downstream model receives one tokenization in both training and inference consistently. We consider that this consistency improves the performance.

3.2 Machine Translation

Settings For experiments on the machine translation task, we employ IWSLT and WMT corpora on eight language pairs. We pre-tokenized all the datasets except for the Chinese corpus with Moses Tokenizer⁹, and we used jieba¹⁰ for the Chinese corpus. We evaluate the performance of each method with detokenized BLEU with SacreBLEU (Post, 2018) after detokenization.

As a recent tokenizer for machine translation, we compare the proposed method with DPE (He et al., 2020), which tokenizes a target sentence, considering the source tokenization, in addition to SentencePiece. We employed the official implementation of DPE¹¹ and train the DPE model using SentencePiece tokenization. In the same as text

⁸We limit the size of the vocabulary to the half size of the initial vocabulary as well as the training of OpTok.

⁹<https://github.com/moses-smt/mosesdecoder>

¹⁰<https://github.com/fxsjy/jieba>

¹¹<https://github.com/xlhex/dpe>

	Encoder	SP	SP+R	SP+R	Ours	SP+R	Ours
	Decoder	SP	SP+R	DPE	SP+R	Ours	Ours
IWSLT14	De-En	33.79	35.03	35.02	34.90	35.78	35.13
	En-De	28.09	29.13	29.39	29.56	29.57	29.30
IWSLT15	Vi-En	28.70	28.78	28.85	29.34	29.69	29.44
	En-Vi	30.87	31.60	31.63	31.41	31.74	31.70
	Zh-En	20.44	21.17	21.38	21.63	21.65	21.89
	En-Zh	14.40	15.25	15.21	15.45	15.59	15.31
IWSLT17	Ar-En	29.23	29.39	29.37	29.48	30.04	29.78
	En-Ar	15.45	17.75	17.83	18.49	18.18	18.21
	Fr-En	37.87	38.43	38.52	38.82	38.68	38.58
	En-Fr	37.95	39.83	39.90	40.01	40.08	39.68
WMT09	Hu-En	14.84	15.51	15.75	15.73	15.74	15.60
	En-Hu	11.02	12.14	12.30	12.30	12.37	12.33
WMT14	De-En	31.46	31.89	31.97	32.19	31.98	31.90
	En-De	27.10	27.41	27.49	27.62	27.52	27.44
WMT16	Ro-En	29.10	31.79	31.80	31.80	31.83	31.72
	En-Ro	21.78	24.05	24.29	24.36	24.53	24.03

Table 2: Results of experiments on machine translation task using IWSLT and WMT corpus (BLEU). We show the tokenization method for Encoder and Decoder. SP and R mean SentencePiece and subword regularization, respectively. The highest scores are highlighted in bold.

classification, we used subword regularization as a strong baseline.

For the downstream model, we used Transformer (Vaswani et al., 2017) implemented in Fairseq (Ott et al., 2019). For the IWSLT dataset, we used the small Transformer, and we created the initial vocabulary using SentencePiece with a 16K size of the vocabulary for each language. For the WMT dataset, we employed Transformer (base), and the size of the vocabulary is 32K. Similar to the case of text classification tasks, we initialized our NULM with the result of SentencePiece. The hyperparameters for subword regularization are $\alpha = 0.2$ for IWSLT, $\alpha = 0.5$ for WMT, and $k = \infty$ for both datasets. The number of tokenizations for the training of the proposed method is $N = 8$ for ISWLT and $N = 3$ for WMT.

In the training of NMT with DPE, we applied subword regularization for the source side language, similar to He et al. (2020). For the proposed method, we prepared three configurations: used our method only for a source side language, only for a target side language, and for both side languages.

Results Table 2 details the performance of each configuration. This table indicates that the system employing our approach achieves the best performance in most datasets. The setting where the proposed method is used only for the decoder side succeeds on many datasets. In contrast, when we use our method for both sides, the performance degrades. These results imply that it is challenging to optimize the tokenization of source and target languages simultaneously, and it can degrade the

performance. We discuss the simultaneous optimization of source and target languages on NMT in Section 5.1.

4 Tokenization as Post-processing

Settings As described in Section 1, our proposed method can be applied as post-processing to an already trained model. In this section, we evaluate the effectiveness of optimizing our tokenizer for the trained model. Concretely, we trained the NULM with \mathcal{L}_s in Eq.(3) without updating the parameters of the downstream model.

We conducted experiments on text classification (Sentiment) and machine translation (IWSLT15) tasks. We trained the downstream models used in Section 3 with subword regularization (Kudo, 2018). We trained the models with 30 epochs for text classification and 100 epochs for machine translation. After the training, we trained only our tokenizer with five epochs using the loss values computed by the trained models. Moreover, for text classification, we trained OpTok in the same manner as our proposed method as a baseline.

Results Table 3 details the performance of each method. This table indicates that the proposed method also improves the performance from the base model trained with subword regularization. The proposed method outperforms OpTok on two datasets of text classification. Moreover, the proposed method increases the BLEU scores consistently in machine translation. These results show that the proposed method is useful to improve the performance of the downstream model even if we

	SP+R	OpTok	Ours
<i>Sentiment</i> (F1)			
Weibo(Zh)	92.69	93.08	92.99
Twitter(Ja)	85.88	86.23	86.28
Twitter(En)	77.21	77.41	77.77
<i>IWSLT15</i> (BLEU)			
Vi-En	28.82	-	28.91
En-Vi	30.48	-	30.60
Zh-En	21.55	-	21.82
En-Zh	14.57	-	14.83

Table 3: Improvement in the performance by optimizing tokenization as post-processing by OpTok and our method. SP+R indicates SentencePiece with subword regularization. The highest scores are highlighted in bold.

use a sufficiently trained model as the downstream model. In other words, since we do not necessarily require training for the proposed method with the downstream model from scratch, our proposed method can be applied to various situations such as the combination with a pre-trained model.

5 Discussion

5.1 Learning Both Encoder and Decoder

The results of the machine translation task (Section 3.2) reveal that the performance decreases when we incorporate our method into both the encoder and the decoder sides. We consider that the cause of this decrease to be the gap in the tokenization strategy between the source and target languages. In this section, we attempt to make it stable to train our method on both the encoder and decoder sides simultaneously with three possible strategies.

Enc→Dec: We train only the encoder-side NULM in the first 50 epochs, with the decoder-side NULM being frozen; then, we train the decoder-side NULM in the last 50 epochs, with the encoder NULM being frozen.

Dec→Enc: We train our method with the reversed version of the strategy Enc→Dec strategy.

Random: We randomly update either of the NULM on the encoder or that on the decoder sides with at a 0.5 ratio in each mini-batch training.

Table 4 presents the results of the experiments. These results indicate that the Enc→Dec strategy contributes to improving the performance of the simultaneous learning of tokenization on both sides. In particular, the scores of Vi-En, En-Vi, and Zh-En surpass the best scores reported in Table 2, indicating that the Enc→Dec strategy is effective for the training of our method. In contrast, the Dec→Enc strategy decreases the performance in many language pairs. The performance obtained us-

	Both	Enc→Dec	Dec→Enc	Random
Vi-En	29.44	30.22	29.47	29.37
En-Vi	31.70	31.78	31.33	31.70
Zh-En	21.89	21.99	21.82	21.66
En-Zh	15.31	15.54	14.88	15.14

Table 4: Performance of NMT on the IWSLT15 datasets with three strategies for the simultaneous training of our method. The scores of Both are taken from the Ours-Ours column in Table 2. The highest scores are highlighted in bold.

ing the Random strategy is slightly lower than that obtained using the original method (Both). From these results, we can conclude that it is effective for the machine translation task to learn the tokenization of each side step-by-step, specifically, from the encoder-side to the decoder-side, instead of optimizing both sides simultaneously.

5.2 Analysis of Tokenization

Optimized Tokenization In this section, we analyze the tokenization obtained using the proposed method on a machine translation task. Table 5 presents the comparison of tokenization among SentencePiece, DPE, and our method. We utilized the IWSLT15 Zh-En corpus for this comparison and tokenized English side sentences using each method. For our method, we only optimized the English side tokenization.

Table 5a presents a comparison of the tokenization on the source side between SentencePiece and the proposed method. Our method splits words into smaller segments than SentencePiece, which is the initial tokenization of our method. For example, our method cuts off the suffix from a stem word, such as splitting “don” into “do-n,” “have” into “hav-e,” and “hours” into “hour-s.”

Table 5b presents a comparison of the tokenization on the target side between SentencePiece, DPE, and the proposed method. Compared to the tokenization on the source side, our method does not split words into tiny units on the target side. The proposed method exhibits the same tendency of tokenization as DPE, such as splitting the past-suffix “-ed.” However, the DPE tokenization contains smaller units than our tokenization; an example of this is the difference in the tokenization for “away.”

Tokenization Granularity To compare the granularities of each tokenizer, we confirm the number of tokens in the corpus tokenized by each method. Table 6 presents the ratio of the number of tokens in the training corpus between the initial tokenization (SentencePiece) and the optimized tokeniza-

SP	Student s don ’ t have long hours of learning .
Ours	Student s do n ’ t hav e long hour s of learning .
TGT	学生 在 校 学 习 时 间 不 长 。

(a) Tokenization difference for the source language (En-Zh)

SRC	引 力 与 其 它 力 分 隔 开 来
SP	Gra vity separate d away from the other force s .
DPE	Gra vity separat ed a way from the other force s .
Ours	Gra vity separat ed away from the other force s .

(b) Tokenization difference for the target language (Zh-En)

Table 5: Comparison of English tokenization on Zh-En pairs using SentencePiece (SP), DPE, and our method. SRC and TGT indicate the tokenization of source and target side, respectively. The different tokenization is highlighted in bold.

tion (DPE and the proposed method). In the table, a value greater than 1 indicates an increase in the number of tokens compared to SentencePiece.

The results reveal that the number of tokens in the proposed method increases for the source side tokenization, which means that our method tokenizes a source corpus into small units by splitting morphemes, as shown in Table 5a.

For the tokenization of the target side, the ratio of the number of tokens for the proposed method is slightly smaller than the initial tokenization, other than for the En-Zh pair. We consider that our method seeks appropriate tokenization to aid in the decoding process while maintaining the granularity of the initial tokenization. With respect to the translation of the En-Zh pair, our method splits a Chinese sentence into smaller tokens. Chinese characters contain much more information than English characters, and the number of Chinese tokens in a sentence is smaller than that of English. We consider that this difference causes the increased tokens on the target side to use the same granularity as the source English corpus.

Compared with the tokenization by the proposed method, the number of tokens for the DPE varies for each language pair. DPE tokenization is more flexible than our method because DPE employs the Transformer and a special decoding algorithm for tokenization, whereas we simply use a unigram language model and the Viterbi algorithm. In addition, DPE tokenizes the target sentence by directly considering the source tokenization by inputting a source sentence to the Transformer. In contrast, we use the target side NULM trained with the both side information to find the target side tokenization. Although our tokenization flexibility is limited, our method improves the performance on NMT tasks, as demonstrated in the experimental results.

Encoder	Ours	SP+R	SP+R
Decoder	SP+R	Ours	DPE
<i>IWSLT14</i>			
De-En	2.5353	0.9992	1.0439
En-De	1.3809	0.9996	0.9923
<i>IWSLT15</i>			
Vi-En	1.5320	0.9993	1.0428
En-Vi	1.4650	0.9999	0.9923
Zh-En	1.5175	0.9994	0.9907
En-Zh	1.3516	1.4713	1.0346
<i>IWSLT17</i>			
Ar-En	2.5350	0.9997	0.9952
En-Ar	1.4765	0.9994	0.9945
Fr-En	1.7194	0.9996	1.0001
En-Fr	1.5996	0.9997	0.9935

Table 6: Ratio of the number of tokens between initial tokenization (SentencePiece) and optimized tokenization (DPE and Ours) on the IWSLT corpora. SP+R denotes SentencePiece with subword regularization.

5.3 Effect of Hyperparameter N

The proposed method updates the NULM using N -best tokenized candidates. In this section, we confirm the effect of the hyperparameter N on the performance for the downstream tasks.

We conducted experiments on text classification and machine translation with different N , as mentioned in Section 3. Figures 3 and 4 show the results respectively. In these figures, we illustrate the difference from the performance of the model with the settings used in Section 3, i.e., $N = 3$ for text classification and $N = 8$ for machine translation.

For the text classification task, we confirm the effect of N on the sentiment analysis datasets. In Figure 3, we can observe that the number of N does not have a strong effect on the performance of the proposed method. In addition, the larger N leads to slightly better performance for the Japanese and English datasets. In contrast, the performance for the Chinese dataset decreases with a large N . We consider that this occurs because a Chinese sentence has more tokenization candidates than the others, and the optimization of tokenization becomes unstable with larger N .

Compared to the existing OpTok method (Hiraoka et al., 2020) the proposed method is robust to large N . As described in Section 3, our method avoids the gap between training and inference in terms of the weighting strategy. Because the proposed method uses one sampled tokenization to train the downstream model, the number of N does not affect the text classification performance. The experiment with various N verifies that our method is superior to Hiraoka et al. (2020).

For the machine translation task, we conduct ex-

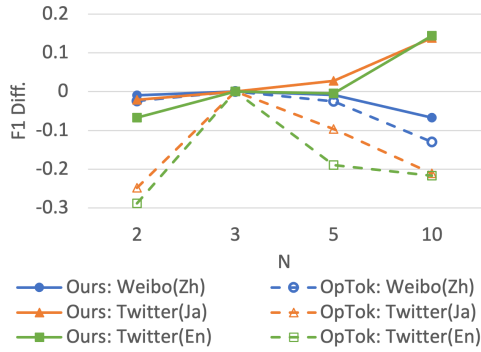


Figure 3: Difference in performance on text classification tasks against N .




Figure 4: Difference in performance on machine translation tasks (Vi-En) against N .

periments on the Vi-En pair of IWSLT15. Figure 4 illustrates that the number of N does not have a strong effect on the performance when we use the proposed method solely for the target side (SP-OURS). When we incorporate our method to the source side (OURS-SP), the performance increases with a large N . We consider that the proposed method is able to seek appropriate tokenization from the large search space when we set a large number as N because the neural encoder of NMT allows various tokenizations for its input. When we use our method for both the encoder and the decoder (OURS-OURS), the performance decreases slightly with higher N . We consider that optimization of the tokenization of both sides with a large N becomes unstable because tokenization on the source side varies vastly during training.

6 Related Work

Many researchers have tackled the problem of optimizing tokenization, especially in the machine translation field. For statistical machine translation, Nießen and Ney (2004) and Goldwater and McClosky (2005) attempted to obtain good tokenization using hand-crafted linguistic information. Some studies explored appropriate tokenization using alignment information between the source

and target languages (Xu et al., 2008; Chung and Gildea, 2009; Nguyen et al., 2010).

Recent studies have attempted to obtain appropriate tokenization for the downstream task using neural networks. Gowda and May (2020) analysed the optimal granularity of tokenization on NMT. Salesky et al. (2020) proposed Incremental-BPE, which automatically explores the appropriate granularity of BPE tokenization. They stopped the merge operation of the BPE depending on the loss on a validation split. He et al. (2020) proposed DPE, which obtains the appropriate tokenization of a target language depending on the tokenization of the source side on the NMT. Our method is different from DPE in that our method can optimize tokenization considering the parameters of the downstream model. Moreover, our method can be applied to both the source and target languages of the machine translation task. Hiraoka et al. (2020) proposed OpTok, which optimizes the tokenizer and the downstream model on text classification simultaneously. We extend their idea to be applicable to any downstream task, including machine translation. Moreover, our method uses a different training strategy to that used in OpTok. We split the training loss for the downstream loss and tokenization loss, as mentioned in Section 2, and the experimental results demonstrate that our strategy is superior to OpTok.

We employ subword regularization to train the downstream model. Kudo (2018) proposed training a model by sampling tokenization with a unigram language model, and Provilkov et al. (2020) modified this idea to use the BPE (Sennrich et al., 2016) process to yield various tokenizations. Hiraoka et al. (2019) applied subword regularization for text classification tasks.

7 Conclusion

We propose a novel method for optimizing tokenization by considering downstream tasks, such as a training corpus and a downstream model. Our method is the first approach to explore appropriate tokenization for any downstream task. Experimental results demonstrate that the proposed method achieves higher performance than existing systems with respect to text classification and machine translation tasks. Because the proposed method is applicable to any architecture using loss for its optimization, we expect our method to improve the performance for other NLP tasks.

Acknowledgments

These research results were obtained from the commissioned research by National Institute of Information and Communications Technology (NICT), Japan.

References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Deng Cai, Hai Zhao, Zhisong Zhang, Yuan Xin, Yongjian Wu, and Feiyue Huang. 2017. [Fast and accurate neural word segmentation for Chinese](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 608–615, Vancouver, Canada. Association for Computational Linguistics.
- Pi-Chuan Chang, Michel Galley, and Christopher D. Manning. 2008. [Optimizing Chinese word segmentation for machine translation performance](#). In *Proceedings of the Third Workshop on Statistical Machine Translation*, pages 224–232, Columbus, Ohio. Association for Computational Linguistics.
- Tagyoung Chung and Daniel Gildea. 2009. [Unsupervised tokenization for machine translation](#). In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 718–726, Singapore. Association for Computational Linguistics.
- Miguel Domingo, Mercedes García-Martínez, Alexandre Helle, Francisco Casacuberta, and Manuel Heranz. 2018. [How much does tokenization affect neural machine translation?](#) *CoRR*, abs/1812.08621.
- Sharon Goldwater, Thomas L. Griffiths, and Mark Johnson. 2006. [Contextual dependencies in unsupervised word segmentation](#). In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 673–680, Sydney, Australia. Association for Computational Linguistics.
- Sharon Goldwater, Thomas L. Griffiths, and Mark Johnson. 2009. A bayesian framework for word segmentation: Exploring the effects of context. *Cognition*, 112(1):21–54.
- Sharon Goldwater and David McClosky. 2005. [Improving statistical MT through morphological analysis](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 676–683, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- Thamme Gowda and Jonathan May. 2020. [Finding the optimal vocabulary size for neural machine translation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3955–3964, Online. Association for Computational Linguistics.
- Ruining He and Julian McAuley. 2016. [Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering](#). In *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, page 507–517, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Xuanli He, Gholamreza Haffari, and Mohammad Norouzi. 2020. [Dynamic programming encoding for subword segmentation in neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3042–3051, Online. Association for Computational Linguistics.
- Tatsuya Hiraoka, Hiroyuki Shindo, and Yuji Matsumoto. 2019. [Stochastic tokenization with a language model for neural text classification](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1620–1629, Florence, Italy. Association for Computational Linguistics.
- Tatsuya Hiraoka, Sho Takase, Kei Uchiumi, Atsushi Keyaki, and Naoaki Okazaki. 2020. [Optimizing word segmentation for downstream task](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1341–1351, Online. Association for Computational Linguistics.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Taku Kudo. 2006. Mecab: Yet another part-of-speech and morphological analyzer. <http://taku910.github.io/mecab/>.
- 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.
- Daichi Mochihashi, Takeshi Yamada, and Naonori Ueda. 2009. [Bayesian unsupervised word segmentation with nested Pitman-Yor language modeling](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 100–108, Suntec, Singapore. Association for Computational Linguistics.
- Hajime Morita, Daisuke Kawahara, and Sadao Kurohashi. 2015. [Morphological analysis for unsegmented languages using recurrent neural network language model](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2292–2297, Lisbon, Portugal. Association for Computational Linguistics.
- Masaaki Nagata. 1994. [A stochastic Japanese morphological analyzer using a forward-DP backward-A* n-best search algorithm](#). In *COLING 1994 Volume 1: The 15th International Conference on Computational Linguistics*.
- ThuyLinh Nguyen, Stephan Vogel, and Noah A. Smith. 2010. [Nonparametric word segmentation for machine translation](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*, pages 815–823, Beijing, China. Coling 2010 Organizing Committee.
- Sonja Nießen and Hermann Ney. 2004. [Statistical machine translation with scarce resources using morpho-syntactic information](#). *Computational Linguistics*, 30(2):181–204.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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.
- 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.
- Rakuten, Inc. 2014. Rakuten dataset. Informatics Research Data Repository, National Institute of informatics. (dataset).
- Elizabeth Salesky, Andrew Runge, Alex Coda, Jan Niehues, and Graham Neubig. 2020. [Optimizing segmentation granularity for neural machine translation](#). *Machine Translation*, pages 1–19.
- Steven L Scott. 2002. Bayesian methods for hidden markov models: Recursive computing in the 21st century. *Journal of the American Statistical Association*, 97(457):337–351.
- 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.
- Kazuma Takaoka, Sorami Hisamoto, Noriko Kawahara, Miho Sakamoto, Yoshitaka Uchida, and Yuji Matsumoto. 2018. [Sudachi: a Japanese tokenizer for business](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Arseny Tolmachev, Daisuke Kawahara, and Sadao Kurohashi. 2018. [Juman++: A morphological analysis toolkit for scriptio continua](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 54–59, Brussels, Belgium. 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. *Advances in neural information processing systems*, 30:5998–6008.
- Andrew Viterbi. 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE transactions on Information Theory*, 13(2):260–269.
- Jia Xu, Jianfeng Gao, Kristina Toutanova, and Hermann Ney. 2008. [Bayesian semi-supervised Chinese word segmentation for statistical machine translation](#). In *Proceedings of the 22nd International Conference on Computational Linguistics (Coling 2008)*, pages 1017–1024, Manchester, UK. Coling 2008 Organizing Committee.
- Jie Yang, Yue Zhang, and Fei Dong. 2017. [Neural word segmentation with rich pretraining](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 839–849, Vancouver, Canada. Association for Computational Linguistics.

Jie Yang, Yue Zhang, and Shuailong Liang. 2019. [Sub-word encoding in lattice LSTM for Chinese word segmentation](#). 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 2720–2725, Minneapolis, Minnesota. Association for Computational Linguistics.

Yongfeng Zhang, Min Zhang, Yi Zhang, Guokun Lai, Yiqun Liu, Honghui Zhang, and Shaoping Ma. 2015. Daily-aware personalized recommendation based on feature-level time series analysis. In *Proceedings of the 24th international conference on world wide web*, pages 1373–1383.

Method	Tokenization
SentencePiece	This didn't have full seasons . I expected more than 6 episodes on this because I liked this series .
<i>Genre (Gold: Movies_and_TV)</i>	
OpTok	This didn't have full seasons . I expected more than 6 episode s on this because I liked this series .
Ours	This didn't have full season s . I expected more than 6 episode s on this because I liked this series .
<i>Rating (Gold: 1)</i>	
OpTok	This didn't have full seasons . I expected mor e than 6 episodes on this because I like d this series .
Ours	This didn't have full seasons . I expected more than 6 episodes on this because I liked this series .

Table 7: Tokenization examples on the text classification task. Bold highlights the difference of tokenization among methods.

A Detailed Experimental Settings

For all the experiments, we implement the proposed method using PyTorch, and we run all the experiments on NVIDIA Tesla V100 (16 GiB). For the initialization of the NULM, we terminate the pretraining when the loss is less than 1×10^{-7} or the training epoch achieves the maximum number (100,000).

A.1 Text Classification

In Section 3.1, we use two new datasets in addition to the datasets used in the existing research (Hiraoka et al., 2020). We prepare both datasets in the same manner as that used for Genre(En) and Rating(En).

For Genre(Zh) and Rating(Zh), we use 13 genres that contain a sufficient number of reviews and sample 30,000 reviews from each genre, balancing the number of ratings and limiting the number of characters in the review to 100. The dataset contains 390,000 reviews, and we split it at a ratio of 8:1:1 for training, validation, and testing. Each review has ratings from 1 to 5 attached by reviewers, and we use the sampled dataset for a rating prediction task and a genre prediction task.

For Genre(Ja) and Rating(Ja), we use 21 genres that contain a sufficient number of reviews and sample 5,000 reviews from each genre and rate, limiting the number of characters in the review to 100. The dataset contains 525,000 reviews, and we split it at a ratio of 8:1:1 for training, validation, and testing. Each review has ratings from 1 to 5 attached by reviewers, and we use them for rating and genre prediction tasks.

We conduct experiments on text classification tasks under the same settings as those used in the existing literature (Hiraoka et al., 2020). Thus, we conduct a text classification with BiLSTM encoders whose hidden size is 256. The size of the word embedding is 64, and we set the batch size to 256 and the maximum training epoch to 20. The pretrained word embeddings are frozen in the train-

ing of text classification, and both NULM and the downstream model share word embeddings.

A.2 Machine Translation

For experiments on machine translation, we do not freeze the word embeddings. We empirically find that the training becomes unstable when word embeddings are shared between NULM and the downstream model without freezing. Therefore, we prepare different word embeddings for NULM and the downstream model, called the Transformer. We set the word embedding size to 64 for the NULM. We make the mini-batch by specifying the number of maximum tokens, and we set it to 1,000. The maximum number of training epochs is 100 for all the experiments, and we average the parameters of the last 10 epochs for evaluation.

B Tokenization for Text Classification

We present examples for the tokenization on text classification tasks, Genre/Rating(En), in Table 7. As both Genre and Rating datasets are created from the same review corpus, we can confirm whether each method can tokenize a sentence depending on the task, which might be a genre prediction or a rating prediction. We compare the tokenization by SentencePiece, OpTok, and our method.

The tendency for tokenization by our method is similar to that by OpTok because our method is based on OpTok. In the example, both OpTok and our method split a suffix “s” from “episodes” only on the genre prediction task. This example implies that both methods yield tokenization that includes task-specific word such as “episode” for the “Movies_and_TV” genre. In addition, our method tokenizes “seasons” into “season-s,” which is also related to the movie genre.

With respect to the rating prediction, our method does not split “liked” into “like,” which might be helpful for predicting ratings, whereas OpTok does. We consider that our method uses the original word to distinguish the verb from the adjective/adverb “like.”