

SubRegWeigh: Effective and Efficient Annotation Weighing with Subword Regularization

Kohei Tsuji¹, Tatsuya Hiraoka², Yuchang Cheng^{1,3}, Tomoya Iwakura^{1,3},

¹NAIST, ²MBZUAI, ³Fujitsu Ltd.
tsuji.kohei.tl1@naist.ac.jp,
tatsuya.hiraoka@mbzuai.ac.ae,
{cheng.yuchang, iwakura.tomoya}@fujitsu.com,

Abstract

NLP datasets may still contain annotation errors, even when they are manually annotated. Researchers have attempted to develop methods to automatically reduce the adverse effect of errors in datasets. However, existing methods are time-consuming because they require many trained models to detect errors. This paper proposes a time-saving method that utilizes a tokenization technique called subword regularization to simulate multiple error detection models for detecting errors. Our proposed method, SubRegWeigh, can perform annotation weighting four to five times faster than the existing method. Additionally, SubRegWeigh improved performance in document classification and named entity recognition tasks. In experiments with pseudo-incorrect labels, SubRegWeigh clearly identifies pseudo-incorrect labels as annotation errors. Our code is available at <https://github.com/4ldk/SubRegWeigh>.

1 Introduction

NLP datasets usually consist of raw texts and annotation labels. Various tasks exploit the pairs of texts and labels for training and evaluating models. To achieve higher performance in NLP tasks, the models should be trained or fine-tuned with a sophisticated training dataset without annotation errors.

However, some popular datasets such as CoNLL-2003 contain annotation errors (Wang et al., 2019; Reiss et al., 2020). When datasets include errors, the performance is degraded by training models from incorrect training datasets, and models are incorrectly evaluated by errors of test datasets.

Recent studies have explored automated annotation using generative AI (Goel et al., 2023; Bogdanov et al., 2024) and the collaboration of humans and AI (Naraki et al., 2024). Even with these methods, it is difficult to avoid annotation errors. Therefore sophisticated **weighing** methods to auto-

matically detect the annotation errors and reduce their negative effect have been expected.

Such a method to weigh annotation errors is recently studied in the NER field. Wang et al. (2019) proposed CrossWeigh, which detects annotation errors in the dataset and adjusts their learning priority by weighting loss values so that the training is not affected by such annotation errors. However, there are shortcomings in its computational efficiency, especially in the recent NLP trends with the pre-trained large language models. Reducing the computational cost contributes to not only speeding up the development of NLP but also Green AI (Schwartz et al., 2019). Furthermore, we cannot utilize CrossWeigh for NLP tasks other than NER because it is specially designed for NER. Therefore, we are required to develop the new weighing methods that can be widely used for various NLP tasks.

In this study, we propose an efficient and effective method of annotation weighing, **SubRegWeigh**. Figure 1 shows an overview of the proposed method. SubRegWeigh evaluates the reliability of annotation for a single sample by feeding various input styles to a model and comparing its outputs. Inspired by Takase et al. (2022), we obtain such various inputs with subword regularization, which yields various tokenization candidates. In other words, SubRegWeigh measures annotation reliability to detect errors by observing the consistency of multiple outputs generated with subword regularization. SubRegWeigh can also be regarded as a method to simulate multiple models for annotation evaluation, which allows for faster weighing of annotation errors by using a single model instead of multiple models required by CrossWeigh. Because the weighing processing of SubRegWeigh is not ad hoc for NER, it is applicable to wide NLP tasks.

We conducted experiments on two NLP tasks: NER and text classification. The experimental results of the CoNLL-2003 demonstrate that SubReg-

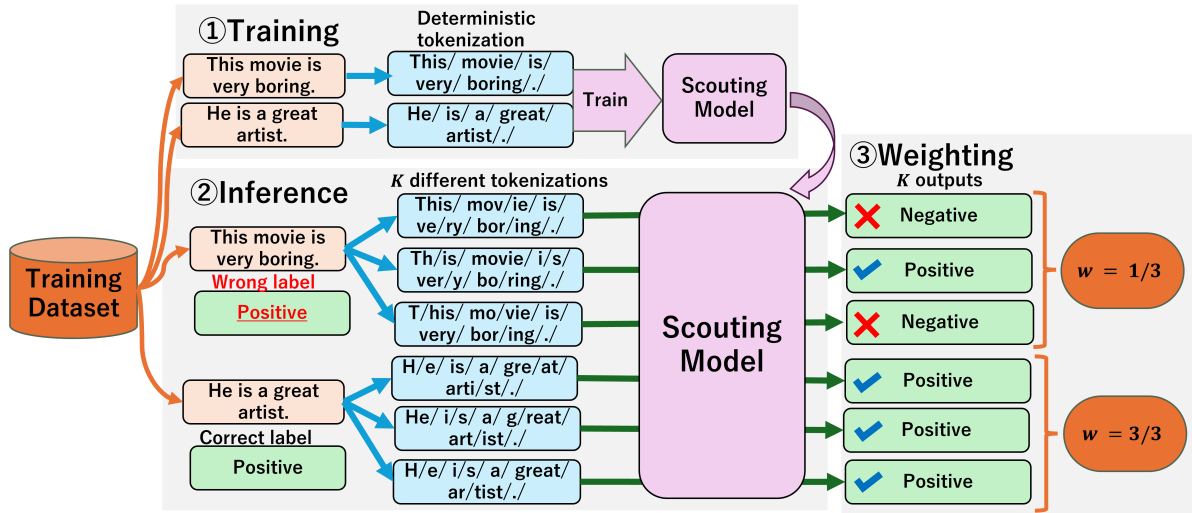


Figure 1: The overview of the three steps of SubRegWeigh with the number of tokenization candidates in the inference $K = 3$ for the classification task whether the input sentence is positive or negative: 1) the training step of the scouting model, 2) the inference step to examine the appropriateness of the training label, and 3) the weighting step to calculate the weight of each training sample according to the appropriateness of their labels. Check marks indicate inferences that are the same as the original label and cross marks indicate inferences that differ from the original label. The calculated weights are used to train the final model.

Weigh can perform the annotation weighing four to five times faster than CrossWeigh. Furthermore, SubRegWeigh contributes to the performance improvement compared to the case using CrossWeigh in some datasets. Especially, our proposed method achieved SoTA on CoNLL-CW, a test dataset constructed by manually correcting annotation errors in the CoNLL-2003 test split. Experiments with pseudo-incorrect labels also showed that the proposed method can recognize annotation errors well.

2 Related Work

2.1 Annotation Error Detection

Our work is in the line of CrossWeigh (Wang et al., 2019), which is an annotation weighing method for NER. CrossWeigh detects annotation errors by performing K -fold cross-validation on the training dataset T times and weights the loss value according to the number of wrong predictions, which degrades the effect of training samples with annotation errors. This method requires large computational costs because it trains $K \times T$ models for detecting annotation errors. Our work is different from CrossWeigh in that we use multiple tokenization candidates to obtain various outputs from one model. This can remarkably reduce the time for detecting annotation errors because we train the model only once and use the output from pseudo-models by inputting various tokenization

candidates. Furthermore, freed from K -fold cross-validation, our method can use the entire training dataset to make the model, which contributes to the more accurate detection of annotation errors.

For the case where we can prepare a small dataset with absolutely correct annotations, some methods have been studied to correct labels directly (Yang et al., 2018; Mayhew et al., 2019; Ding et al., 2023). In other lines, Zhou and Chen (2021) proposed a method to infer the correctness of the labels during training, which can be applied in conjunction with our method.

For detecting annotation errors in datasets beyond NER, annotation correction methods with distant supervision (Mintz et al., 2009) have been studied in the relation extraction tasks (Huang and Du, 2019; Qin et al., 2018). However, these methods require an additional clean dataset or are not always applicable. Furthermore, methods for correcting POS tagging tasks have also been studied (Nakagawa and Matsumoto, 2002; Helgadóttir et al., 2014), but these do not apply to tasks other than POS tagging. The analyzing method of the models' memorization (Maini et al., 2023) could be used to detect training samples with ambiguous labels.

2.2 Subword Regularization

The proposed method exploits the technique of tokenization for efficient annotation error detection.

Tokenization (Song et al., 2021; Sennrich et al., 2016; Kudo, 2018) is widely used in recent NLP. This technique reduces the number of unknown and low-frequency words and limits the number of dictionaries by treating words as a sequence of finer subwords. However, because tokenization is deterministically performed on training data, the subword sequences of inference data may not always match the sequences learned during training.

Subword regularization (Kudo, 2018; Provilkov et al., 2020; Hiraoka, 2022) alleviates this problem of deterministic tokenization. This technique generates multiple tokenization candidates for a single text. While subword regularization is typically used during the training step to learn from various tokenizations, some studies use it during inference. Takase et al. (2022) uses subword regularization during inference to ensemble the various outputs of a single model by inputting some different tokenization candidates to the model. Wang et al. (2021a) proposed a method to reduce the effect of tokenization difference. However, there has been no research using subword regularization during inference to reduce noise in the training data.

Subword regularization using a unigram language model (Kudo, 2018) is a method to sample tokenization candidates using the unigram probability of tokens. For the cases where the model does not employ a tokenizer with the unigram language model like Byte Pair Encoding (BPE) (Sennrich et al., 2016) used in RoBERTa (Liu et al., 2019) or WordPiece (Song et al., 2021) used in BERT (Devlin et al., 2019), we can use BPE-Dropout (Provilkov et al., 2020) and MaxMatch-Dropout (Hiraoka, 2022) for the subword regularization, respectively. To obtain various tokenization candidates, BPE-Dropout randomly avoids merges during the combination process and MaxMatch-Dropout randomly rejects the longest match candidates for each word.

3 Proposed Method: SubRegWeigh

The purpose of the proposed method, SubRegWeigh, is to detect the incorrectly annotated samples in the training dataset and reduce the effect of such a sample with annotation errors during the training of the final model by weighting the loss values. The process of SubRegWeigh is composed of the following three steps (Figure 1):

1. Training a **scouting model** in a usual manner of NLP tasks with the training data using

deterministic tokenization (§3.1).

2. Inferring the labels for the training samples with the trained scouting model. Herein, we tokenize the input texts into K different tokenization candidates. Therefore, we can collect K inferred outputs for each tokenization candidate of a single input text (§3.2).
3. Calculating weights for training samples according to the number of correctly answered inputs for each tokenization candidate (§3.3).

The first and second steps detect the annotation errors by scouting the dataset. Then, SubRegWeigh weighs the correctness of annotation labels and assigns the weight for each training sample in the third step. After the above steps, we can train **the final model** with the original training data and the calculated weights for each training sample (§3.4). Compared to CrossWeigh, SubRegWeigh can be used for various NLP tasks beyond NER because the entire steps are not specialized to NER. The following section explains each step in detail.

3.1 Training of the Scouting Model

Inspired by CossWeigh (Wang et al., 2019), we first create the scouting model M (see the left top of Figure 1). This model is used to scout the correctness of the annotation in the training dataset in the next step. Let (X, Y) be a pair of the text and the label corresponding to a single training sample. When the text is composed of I words, $X = x_1, \dots, x_I$. Y is a single value in text classification (e.g., positive and negative labels) or a sequence of labels $Y = y_1, \dots, y_I$ in sequential labeling such as NER.

In this step, we train the scouting model M in the usual manner of training or fine-tuning for general NLP datasets. The single input text X is tokenized into a sequence of subwords $X' = s_1, \dots, s_J$ that is composed of J subwords¹. Then, we train the scouting model M with the tokenized texts and labels. Note that we use the deterministic tokenization here (e.g., the default encoding of the BPE tokenizer or the 1-best tokenization with the Viterbi algorithm for the unigram language model-based tokenizer). In other words, we do not use subword regularization (i.e., stochastic tokenization) for the training of the scouting model M .

¹The numbers of words I and subwords J are not necessarily matched depending on the tokenization methods. In the case of sequential labeling tasks with $I \neq J$, we aligned the subwords and labels following the existing literature (Liu et al., 2019; Yamada et al., 2020).

3.2 Inference with the Scouting Model

In the second step, we use the trained scouting model M to scout the correctness of the annotated label in the training dataset (see the right part of Figure 1). We assume that the trained scouting model M can correctly predict the labels even for the differently tokenized input if the label is correctly annotated. Concretely, the input text X is tokenized into K different tokenization candidates $\{X'_1, \dots, X'_k, \dots, X'_K\}$. Then each tokenization candidate X'_k is fed into the trained model M and we obtain the corresponding prediction \hat{Y}_k ². We then evaluate the output \hat{Y}_k with the gold annotation Y . This process can be regarded as the “mistake-reweighing” process of CrossWeigh with K pseudo-models, which is realized with the K different tokenization candidates for the single input text and the single model (Takase et al., 2022).

We can obtain various tokenization candidates using a technique of subword regularization. When we use BPE or WordPiece, BPE-Dropout and MaxMatch-Dropout can be used, respectively. When one uses a unigram language model-based tokenizer, we can also use N -best tokenization instead of subword regularization.

We consider it important to make a variety of the K tokenization candidates to weigh the correctness of the annotation from diverse aspects. To avoid using similar subword sequences as the inputs for the model M , we select K more different candidates from a large number N of tokenization candidates. We propose three options of the way to select such K tokenization candidates as follows.

Random: We randomly sample $N = K$ different tokenization candidates using subword regularization and use them as inputs for the model M . In other words, we do not select the tokenization candidates from the large number of candidates.

Cos-Sim: We select K tokenization candidates using the cosine similarity calculated with TF-IDF vectors. TF-IDF was calculated with each subword as a term, each subword sequence as a document, and the entire candidate set as a corpus. First, we convert $N (> K)$ tokenization candidates into TF-IDF vectors and select the first candidate with the smallest cosine similarity against the deterministic tokenization X' . Subsequently, we select the candidate that is least similar to both X' and al-

²As well as Y , \hat{Y}_k can be single or multiple labels depending on target tasks

ready selected candidates. We repeatedly select the candidates until the number of candidates reaches K .

K -means: We select K tokenization candidates using the K -means clustering. We apply the clustering to $N (> K)$ candidates vectorized with TF-IDF. We then choose K candidates at the nearest point to the centroids of each cluster.

Table 6 in Appendix A shows the tokenization examples by each selection method.

3.3 Weighting for Training Samples

To reduce the effect of training samples with annotation errors, we calculate weights for each sample depending on the results of the inference step (see the left bottom of Figure 1). Concretely, we calculate the weight w corresponding to the training sample (X, Y) as the follows:

$$w = \min(w_{\min}, \frac{1}{K} \sum_{k=1}^K c_k), \quad (1)$$

$$c_k = \begin{cases} 1 & \text{if } \hat{Y}_k = Y \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where \hat{Y}_k is the output of the scouting model M when inputting the tokenization candidate X'_k . If all predictions were different from the original label Y , the weights would be 0 and the data would not be used for training. We use a pre-defined minimum weight w_{\min} to avoid the zero weights, following the suggestion in Wang et al. (2019) finding it better to give a small weight to samples with annotation errors than not to use them.

3.4 Training of Final Model

Using the training dataset with weights calculated in the above section, we train the final model M' that is independent of M . Following CrossWeigh (Wang et al., 2019), the loss value L_{weighted} for the training sample (X, Y) with the weight w is defined as the follows:

$$L_{\text{weighted}}(X, Y, w) = wL_{\text{original}}(X, Y), \quad (3)$$

where $L_{\text{original}}(\cdot)$ is the loss function used in the default training.

4 Experiments

We examine the processing time for the annotation error detection and weighting. Besides, we evaluate the performance of NER and text classification when using the training dataset with the weights.

4.1 Dataset

We briefly overview the dataset used in the experiments here. Appendix B also explains in detail.

NER: We used **CoNLL-2003** (Tjong Kim Sang and De Meulder, 2003) with the official split for the training, validation, and test. The label format was changed from IOB1 to IOB2 (Tjong Kim Sang and Veenstra, 1999). In addition, we used **CoNLL-CW** (Wang et al., 2019) and **CoNLL-2020** (Liu and Ritter, 2023) for the test dataset.³

Text Classification: We used **SST-2** (Socher et al., 2013). The official training split was used for the training dataset and the validation split was used for the test dataset.

4.2 Compared Methods

We compared the following five methods including three options of SubRegWeigh explained in §3.2.

For the baselines, we selected **Vanilla** and **CrossWeigh**. **Vanilla** is the final model that is trained on the original training split without any annotation weighing. **CrossWeigh** is the final model that is trained with the dataset weighed by the official code⁴ of the existing work (Wang et al., 2019). Although CrossWeigh is specialized to NER, we forcibly utilize it for the text classification setting without Entity Disjoint (Wang et al., 2019).

We evaluated the three options of our method introduced in §3.2, which are represented as **SubRegWeigh (Random)**, **SubRegWeigh (Cos-Sim)**, and **SubRegWeigh (K -means)**, respectively. We used these proposed methods to weight the training dataset and evaluate the final model trained with this weighted dataset as well as CrossWeigh.

4.3 Model Settings

We employed RoBERTa_{LARGE} as the pre-trained language model for both NER and text classification. Besides LUKE_{LARGE} (Yamada et al., 2020) is used for the NER setting because LUKE is an architecture for entity-related tasks. In this experimental setting, we used the same backbone model for both the scouting model and the final model⁵. For example, when RoBERTa is used in the scouting model, the final model will also use RoBERTa.

³Because both of these dataset names are CoNLL++, we call CoNLL-CW and CoNLL-2020 to distinguish them.

⁴<https://github.com/ZihanWangKi/CrossWeigh>

⁵§5.3 discusses the case where using different pre-trained models between the scouting and final models.

The detailed model settings such as hyperparameters are provided in the Appendix C.

We set the three hyperparameters of CrossWeigh: the number of folds in mistake estimation $K = 10$, the number of iterations of mistake estimation $T = 3$, and the weight scaling factor $\epsilon = 0.7$. BPE-Dropout (Provilkov et al., 2020) was used to obtain the multiple tokenization candidates for SubRegWeigh because both RoBERTa and LUKE employ BPE for their tokenizers. BPE-Dropout has a hyperparameter p , where a higher p results in finer tokenization (i.e., more different from the original tokenization). In the experiments, p was set to 0.1 unless otherwise noted. Additionally, by default, the hyperparameters for SubRegWeigh were set to $N = 500$, $K = 10$, and $w_{\min} = 1/3$.

4.4 Evaluation Metrics

We evaluate the methods from the viewpoints of speed and performance. On the speed side, we measured the time for detecting the annotation errors and generating the weighted data for the training dataset in NER⁶. On the performance side, we measured the F_1 score in NER and the accuracy in text classification on each test dataset for the final models trained on the weighted training data.

4.5 Experimental Results

The results of our experiments are shown in Table 1. We reported the averaged scores and the standard deviations over the five independent runs.

4.5.1 Processing Time

The column “Processing time” shows the time taken to detect the annotation errors and assign weights for each sample in the training dataset in NER. As shown in this column, the proposed method, SubRegWeigh (Random), was maximum of over eight times faster than CrossWeigh. Even the most time-consuming method of SubRegWeigh (K -means) was five times faster with RoBERTa and four times faster with LUKE than CrossWeigh. The method with the random selection is the fastest among the proposed methods because it does not require any of the calculations about TF-IDF, cosine similarity, and K -means clustering. In addition, in the random method, only $K(= N = 10)$ tokenization candidates were generated, whereas the other methods of SubRegWeigh need to select K

⁶We measured the processing time with Xeon Platinum 8167M + NVIDIA V100 for RoBERTa and Xeon Gold 6230R + NVIDIA A100 for LUKE.

	Processing time hh:mm	CoNLL valid F_1	CoNLL test F_1	CoNLL CW F_1	CoNLL 2020 F_1	SST-2 ACC
SoTA	-	-	94.60 [†]	95.88 ^{††}	-	-
RoBERTa _{LARGE}						
Vanilla	-	97.26 \pm 0.12	93.54 \pm 0.28	95.27 \pm 0.24	94.80 \pm 0.22	94.68 \pm 0.12
CrossWeigh	30:55	97.11 \pm 0.11	93.40 \pm 0.21	94.99 \pm 0.19	94.93 \pm 0.25	94.31 \pm 0.09
SubRegWeigh						
Random	3:26	97.30 \pm 0.16	93.51 \pm 0.26	95.24 \pm 0.22	94.52 \pm 0.22	94.61 \pm 0.10
Cos-Sim	4:51	97.27 \pm 0.12	93.44 \pm 0.21	95.17 \pm 0.23	94.91 \pm 0.19	94.75 \pm 0.09
K-means	5:21	97.28 \pm 0.11	93.81 \pm 0.16	95.45 \pm 0.25	94.96 \pm 0.21	94.84 \pm 0.07
LUKE _{LARGE}						
Vanilla	-	96.78 \pm 0.08	94.32 \pm 0.15	95.92 \pm 0.13	95.29 \pm 0.18	-
CrossWeigh	26:19	96.62 \pm 0.08	94.12 \pm 0.19	95.96 \pm 0.12	95.32 \pm 0.18	-
SubRegWeigh						
Random	2:59	96.54 \pm 0.11	94.22 \pm 0.13	95.94 \pm 0.15	95.24 \pm 0.20	-
Cos-Sim	6:19	96.53 \pm 0.10	94.11 \pm 0.16	95.93 \pm 0.13	95.21 \pm 0.25	-
K-means	6:36	96.65 \pm 0.09	94.20 \pm 0.15	96.12 \pm 0.18	95.31 \pm 0.14	-

Table 1: The processing time and the performance of the final models with the weighted dataset. The best performing scores are highlighted in bold. [†] and ^{††} is SoTA score from Wang et al. (2021b) and Zhou and Chen (2021).

candidates from the $N(= 500)$ -sampled tokenization candidates, which results in longer processing time. Although the proposed method of Cos-Sim and K -means takes longer time than Random, it is remarkably faster than the existing method for the automatic annotation weighing, CrossWeigh.

4.5.2 Performance on NER

The performance on CoNLL-CW is the most important in our evaluation because annotation errors in the test split are removed, which indicates that we can compare the methods on the most unpoluted dataset⁷. One can see that SubRegWeigh (K -means) achieves higher performance than the vanilla method on the test dataset of CoNLL-CW. Specifically, SubRegWeigh (K -means) improved the F_1 scores by 0.18 points with RoBERTa and 0.20 points with LUKE compared to each vanilla baseline. This result indicates that the proposed method of annotation weighing contributes to the performance improvement in the clean test dataset. Furthermore, SubRegWeigh (K -means) achieves higher scores than CrossWeigh, which demonstrates that the proposed method is a reasonable alternative to the existing method.

Both SubRegWeigh and CrossWeigh scored lower than Vanilla in the CoNLL valid and test datasets, which contain annotation errors. These results indicate the remarkable negative effect of the annotation errors in the valid and test datasets and the difficulty of evaluating the models appro-

priately.

Among the proposed methods, Random scores lower than K -means in many cases. We consider this because the random selection of tokenization candidates leads to the selection of similar subword sequences, causing a bias in the inference results. Cos-Sim also scored lower than K -means in all cases, which shows that the naive method of selecting tokenization candidates does not contribute to performance improvement. We consider that the K -means clustering can handle the diverse range of subword tokenization candidates, and this leads to the successful result.

4.5.3 Performance on Text Classification

For text classification, SubRegWeigh (K -means) achieved the highest accuracy, and SubRegWeigh (random) had lower accuracy, similar to the results in NER. One reason for the lack of performance improvement with CrossWeigh in SST-2 is that Entity Disjoint is important for CrossWeigh but unavailable in the text classification task.

From the results on the NER and text classification datasets, we conclude that the proposed method is superior to the existing method in terms of processing time and task performance. Especially, we can say SubRegWeigh (K -means) is the best option, which achieves the highest task performance with approximately four to five times faster than CrossWeigh.

⁷The vanilla baseline score of LUKE exceeds the SoTA score (Zhou and Chen, 2021) for CoNLL-CW, which we consider because of the successful hyperparameter setting.

	\bar{w}_{cor}	\bar{w}_{incor}	CoNLL test F_1	CoNLL CW F_1
Vanilla	-	-	84.15	85.49
CrossWeigh	0.8000	0.0019	84.28	85.71
SubRegWeigh				
Random	0.9278	0.0048	83.77	85.04
Cos-Sim	0.8346	0.0042	84.25	85.69
K-means	0.9284	0.0048	84.34	85.76

Table 2: Averaged weights and performance (F_1) for the pseudo-incorrect dataset. \bar{w}_{cor} and \bar{w}_{incor} are the averaged weights for the samples with original labels and ones with the pseudo-incorrect labels, respectively.

5 Discussion

5.1 Pseudo-incorrect Label Test

We verified whether SubRegWeigh can accurately detect annotation errors in the dataset. We evaluate this capability using a modified dataset with some labels flipped to incorrect labels artificially.

Assuming that the original label should be correctly annotated, we replaced 10% of the labels in the CoNLL2003 training dataset with different labels that do not conflict with the IOB2 format (Appendix D). Then, 3,329 sentences including pseudo-incorrect labels, and 5,356 sentences with original annotations are obtained.

We assign weights to this modified dataset using CrossWeigh and SubRegWeigh⁸. Then we compared the averaged weights between the pseudo-incorrect and the original samples to confirm whether annotation error detection is correct. Here, we calculate the weights as $\frac{C}{K}$, where C is the number of inferences that the model predicts the original label, instead of w in Eq. (1). This is because we consider that the minimum weight w_{min} and the weight correction in CrossWeigh could have a negative effect on the fair analysis.

In Table 2, \bar{w}_{incor} is the averaged weights assigned to the samples with pseudo-incorrect labels, where we expect them should be lower as incorrect labels. Similarly, \bar{w}_{cor} is the averaged weights assigned to the original labels, where higher weights should be assigned as the correct labels.

From the results in this table, SubRegWeigh (K -means) assigns the most distinct weights between the replaced and the unreplaced data. Additionally, all methods assign weights approximately 100 times lower to replaced data than unreplaced data. This indicates that all methods can detect errors

⁸The hyperparameter settings are the same as those described in Section §4.3 and Appendix C.

K	Method	Time	CoNLL test	CoNLL CW
500	Random	75:41	93.71	95.26
50	Random	9:15	93.46	95.08
50	Cos-sim	10:52	93.46	95.12
50	K -means	12:02	93.65	95.30
10	Random	3:26	93.51	95.24
10	Cos-Sim	4:51	93.44	95.17
10	K -means	5:21	93.81	95.45

Table 3: Difference in the processing time (hh:mm) and the performance (F_1) against three options of K .

Final Model	Scouting Model	CoNLL test	CoNLL CW
RoBERTa	RoBERTa	93.81	95.45
	LUKE	93.65	95.42
LUKE	RoBERTa	94.24	95.89
	LUKE	94.20	96.12

Table 4: Difference in the performance (F_1) when using different pre-trained models for scouting and final model.

effectively. The lowest average weight for unreplaced data is CrossWeigh. This is because CrossWeigh infers only 3 times while SubRegWeigh infers 10 times. Even one incorrect prediction reduced the weight to 2/3. Among SubRegWeigh methods, Cos-Sim assigns lower weights to sentences without pseudo-incorrect labels, but these do not mean lower performance improvements compared to other methods, especially SubRegWeigh (Random). According to this, if the weights of sentences containing errors are low, it does not matter if the average weight is slightly lower.

In addition to the investigation of weights, we also analyze the obtained performance by training final models with the modified dataset. For the training of the final models, we used the weights calculated in the manner of Eq. (1). The results are shown in the right two columns of Table 2. Because the training data is noisy, the entire scores are worse than the ones in Table 1. However, the total tendency of the performance is similar to the results in §4. From the entire results, we cannot find a clear relationship between the averaged weights and the final performance.

5.2 Numbers of Tokenization Candidates

K is an important hyperparameter affecting both processing time and performance of SubRegWeigh. We investigated the difference in the processing time and the F_1 score of the final model with three options of the number of tokenization candidates K

Text _I Label	CrossWeigh	SubRegWeigh		
		Random	Cos-Sim	K-means
The _O foreign _O ministry _O 's _O Shen _{B-ORG} told _O Reuters _{B-ORG} Television _{I-ORG} in _O an _O interview _O he _O had _O read _O reports _O of _O Tang _{B-PER} 's _O comments _O ...	0.343	0.333	0.333	0.333
<u>NÖTES</u> _O <u>BAYERISCHE</u> _{B-ORG} <u>VEREINSBANK</u> _{I-ORG} <u>IS</u> _O JOINT _O LEAD _O MANAGER _O	0.700	0.500	0.500	0.333
Former _O Surinam _{B-LOC} rebel _O leader _O held _O after _O shooting _O . _O	1.000	0.700	0.700	0.500

Table 5: Examples of weights assigned by each method. Underlined are incorrect or ambiguous labels.

in the NER dataset. We examined the difference in the range of $K = 10, 50, 500$ with RoBERTa. The dropout rate for BPE-Dropout was $p = 0.1$. Since K -means and Cos-Sim are techniques to reduce the variations in subword segmentation, a large value like $K = 500$ makes them almost indistinguishable from the Random method. Therefore, only the Random method was investigated for $K = 500$.

The results are shown in Table 3. Comparing the results between the ones with $K = 10$ and $K = 50$, one can see that K does not have a large effect on the F_1 scores. However, the larger K remarkably increases the processing time. For $K = 500$, the F_1 score improved compared to Random with $K = 10$ and $K = 50$ on the CoNLL and CoNLL-CW datasets. This suggests that using a large number of subword sequences can effectively weigh annotation errors. However, such much large K damages the speed of annotation weighing, making it a trade-off between the speed vs. the performance. While increasing K to 500 in the Random method showed performance improvement, K -means exhibited the highest performance even with smaller K . This indicates that K -means can select sufficiently diverse subword sequences.

5.3 Weighting with Different Backbone

In the experiment shown in §4, we used the same backbone for the scouting model M for annotation weighing and the final model M' . Herein, we investigated the performance difference when using different backbones in NER. Specifically, we examined the effect in F_1 scores when M was RoBERTa and M' was LUKE. Similarly, we examined the case where M was LUKE and M' was RoBERTa. Both models were trained with the hyperparameters explained in Table 7. We used K -means for selecting the tokenization candidates because it showed the best performance in the main results (§4).

The results in Table 4 indicate that, in most cases, the data weighted by the same model tends

to achieve higher f_1 scores. Although we can save time by preparing the weighed dataset with a single model and reusing it to train models with other backbones, this observation suggests that we should prepare the weighed dataset depending on the backbone used for the final models to obtain performance improvement. This suggestion also supports the importance of the fast method of annotation weighing.

5.4 Qualitative Analysis

Several examples of the CoNLL-2003 training dataset weighed by each method are shown in Table 5. All methods successfully assigned low weights for clearly incorrect sentences like the top example in the table. However, for ambiguous labels like the example at the bottom of the table, CrossWeigh tended to assign high weights, whereas SubRegWeigh more frequently assigned low weights. Additionally, we discovered a specific weakness of SubRegWeigh: it tends to assign lower weights to sentences composed entirely of uppercase letters (see the example in the middle of the table). This is because only a few uppercase-only subwords are included in the tokenizer’s vocabulary, and even with a low p , subword regularization significantly changes the tokenization candidates for uppercase-only sentences, leading to incorrect inferences. We believe this issue is not unique to our method but rather a general problem with subword regularization, which we plan to investigate further in future research.

6 Conclusion

We proposed SubRegWeigh, a method for annotation weighing using subword regularization, which offers faster annotation weighing than existing methods. In particular, subword sequence selection using K -means was four to five times faster than CrossWeigh for annotation weighing and contributed to better performance than weighting with

all of the large number of subword sequences.

In addition, the performance dropped in many cases when different models were used for the scouting and final model, indicating the need for comparison including annotation weighing for a better model and the importance of developing an efficient annotation weighing method.

Limitation

In this paper, we proposed a method for annotation weighing using subword regularization. However, since a deep learning model is used for error weighing, the calculated weights are not always guaranteed to be appropriate.

As mentioned in 5.4, if there is not enough variation of the subword sequence, such as uppercase-only sentences, our proposed method often assigns incorrect weights.

In §5.1, we created pseudo-incorrect labels by randomly replacing the original labels with incorrect labels. This error distribution may differ from real-world annotation error distribution.

Ethical Considerations

Experiments presented in this work were performed with existing datasets (Tjong Kim Sang and De Meulder, 2003; Wang et al., 2019; Liu and Ritter, 2023; Socher et al., 2013). These datasets were used to study NER or text classification models, which is consistent with their intended use. Because this study focuses on efficient annotation error detection in the dataset, its potential risks and negative impact on society appear to be minimal.

References

- Sergei Bogdanov, Alexandre Constantin, Timothée Bernard, Benoit Crabbé, and Etienne Bernard. 2024. Nuner: Entity recognition encoder pre-training via llm-annotated data. *arXiv preprint arXiv:2402.15343*.
- Leon Derczynski, Eric Nichols, Marieke van Erp, and Nut Limsopatham. 2017. [Results of the WNUT2017 shared task on novel and emerging entity recognition](#). In *Proceedings of the 3rd Workshop on Noisy User-generated Text*, pages 140–147, Copenhagen, Denmark. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Liangping Ding, Giovanni Colavizza, and Zhixiong Zhang. 2023. [Partial annotation learning for biomedical entity recognition](#). *arXiv preprint arXiv:2305.13120*.
- William B. Dolan and Chris Brockett. 2005. [Automatically constructing a corpus of sentential paraphrases](#). In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Akshay Goel, Almog Gueta, Omry Gilon, Chang Liu, Sofia Erell, Lan Huong Nguyen, Xiaohong Hao, Bolous Jaber, Shashir Reddy, Rupesh Kartha, et al. 2023. [Llms accelerate annotation for medical information extraction](#). In *Machine Learning for Health (ML4H)*, pages 82–100. PMLR.
- Sigrún Helgadóttir, Hrafn Loftsson, and Eiríkur Rögnvaldsson. 2014. [Correcting errors in a new gold standard for tagging Icelandic text](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*, pages 2944–2948, Reykjavik, Iceland. European Language Resources Association (ELRA).
- Tatsuya Hiraoka. 2022. [MaxMatch-dropout: Subword regularization for WordPiece](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4864–4872, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Yuyun Huang and Jinhua Du. 2019. [Self-attention enhanced CNNs and collaborative curriculum learning for distantly supervised relation extraction](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 389–398, Hong Kong, China. Association for Computational Linguistics.
- Taku Kudo. 2018. [Subword regularization: Improving neural network translation models with multiple subword candidates](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Shuheng Liu and Alan Ritter. 2023. [Do CoNLL-2003 named entity taggers still work well in 2023?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8254–8271, Toronto, Canada. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.

- Pratyush Maini, Michael C Mozer, Hanie Sedghi, Zachary C Lipton, J Zico Kolter, and Chiyuan Zhang. 2023. Can neural network memorization be localized? In *International Conference on Machine Learning*.
- Stephen Mayhew, Snigdha Chaturvedi, Chen-Tse Tsai, and Dan Roth. 2019. **Named entity recognition with partially annotated training data**. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 645–655, Hong Kong, China. Association for Computational Linguistics.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. **Distant supervision for relation extraction without labeled data**. 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 1003–1011, Suntec, Singapore. Association for Computational Linguistics.
- Tetsuji Nakagawa and Yuji Matsumoto. 2002. **Detecting errors in corpora using support vector machines**. In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Yuji Naraki, Ryosuke Yamaki, Yoshikazu Ikeda, Takafumi Horie, and Hiroki Naganuma. 2024. Augmenting ner datasets with llms: Towards automated and refined annotation. *arXiv preprint arXiv:2404.01334*.
- 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.
- Pengda Qin, Weiran Xu, and William Yang Wang. 2018. **Robust distant supervision relation extraction via deep reinforcement learning**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2137–2147, Melbourne, Australia. Association for Computational Linguistics.
- Frederick Reiss, Hong Xu, Bryan Cutler, Karthik Muthuraman, and Zachary Eichenberger. 2020. **Identifying incorrect labels in the CoNLL-2003 corpus**. In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 215–226, Online. Association for Computational Linguistics.
- Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2019. **Green AI**. *CoRR*, abs/1907.10597.
- 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.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. **Recursive deep models for semantic compositionality over a sentiment treebank**. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Xinying Song, Alex Salcianu, Yang Song, Dave Dopson, and Denny Zhou. 2021. **Fast WordPiece tokenization**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2089–2103, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Sho Takase, Tatsuya Hiraoka, and Naoaki Okazaki. 2022. **Single model ensemble for subword regularized models in low-resource machine translation**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2536–2541, Dublin, Ireland. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. **Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition**. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Erik F. Tjong Kim Sang and Jorn Veenstra. 1999. **Representing text chunks**. In *Ninth Conference of the European Chapter of the Association for Computational Linguistics*, pages 173–179, Bergen, Norway. Association for Computational Linguistics.
- Xinyi Wang, Sebastian Ruder, and Graham Neubig. 2021a. **Multi-view subword regularization**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 473–482, Online. Association for Computational Linguistics.
- Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021b. **Automated concatenation of embeddings for structured prediction**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2643–2660, Online. Association for Computational Linguistics.
- Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021c. **Improving named entity recognition by external context retrieving and cooperative learning**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1800–1812, Online. Association for Computational Linguistics.
- Zihan Wang, Jingbo Shang, Liyuan Liu, Lihao Lu, Jiacheng Liu, and Jiawei Han. 2019. **CrossWeigh**:

Training named entity tagger from imperfect annotations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5154–5163, Hong Kong, China. Association for Computational Linguistics.

Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yuji Matsumoto. 2020. **LUKE: Deep contextualized entity representations with entity-aware self-attention**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6442–6454, Online. Association for Computational Linguistics.

Yaosheng Yang, Wenliang Chen, Zhenghua Li, Zhengqiu He, and Min Zhang. 2018. **Distantly supervised NER with partial annotation learning and reinforcement learning**. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2159–2169, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Wenxuan Zhou and Muhao Chen. 2021. **Learning from noisy labels for entity-centric information extraction**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5381–5392, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

A Example For Each Selection Method

In §3.2, we use 3 subword sequence selection methods: Random, Cos-Sim, and K-means. We show specific examples of the subword sequences obtained by each selection method in Table 6.

B Detail About Dataset

B.1 NER

The overview in the case of NER is shown in Figure 2.

We use the training split of **CoNLL-2003** (Tjong Kim Sang and De Meulder, 2003) as the target of the annotation weighing. This split is also used to train the final NER models evaluated on the following test datasets. The training split comprises 946 articles (14,041 sentences and 23,499 named entity labels).

We used the validation and test split of **CoNLL-2003** to evaluate the finally trained NER models. The validation split contains 216 articles (3,250 sentences and 5,942 named entity labels) and the test split contains 232 articles (3,453 sentences and 5,648 labels).

In addition to the original CoNLL dataset, we employ the modified version and the recently published version of this dataset for further evaluation.

CoNLL-CW (Wang et al., 2019) is constructed by manually correcting annotation errors in the CoNLL-2003 test split, which includes 231 articles (3,453 sentences and 5648 named entity labels). **CoNLL-2020** (Liu and Ritter, 2023) is constructed with articles from the year 2020 using the same definitions of NER labels as CoNLL-2003, which consists of 131 articles (1,840 sentences and 4,007 named entity labels).

B.2 Text Classification

We used the training and validation split of **SST-2** (Socher et al., 2013). The training split comprises 67,349 sentences. 55.8 % of this split are negative labels, and the rest are positive labels. The validation split comprises 872 sentences. 50.9 % of this split are negative labels, and the rest are positive labels. Test split is not used because the labels have not been published.

C Detail About Model Setting

We employed $\text{RoBERTa}_{\text{LARGE}}$ (Liu et al., 2019) and $\text{LUKE}_{\text{LARGE}}$ (Yamada et al., 2020) for the scouting model M and the final model M' . The scouting and the final model were trained independently. Each model was trained using the hyperparameters in Table 7.

We used the same architectures for the scouting and final models. In other words, the scouting model with $\text{RoBERTa}_{\text{LARGE}}$ is used to weigh the annotation errors that are used for the training of the final model with $\text{RoBERTa}_{\text{LARGE}}$. When the final model is based on $\text{LUKE}_{\text{LARGE}}$, we used the scouting model with $\text{LUKE}_{\text{LARGE}}$.

In NER, for $\text{RoBERTa}_{\text{LARGE}}$, inference was performed with a token-level classification model where the first subword of each word was classified based on BIO tags, and the outputs of the second and subsequent subwords were masked and ignored. For $\text{LUKE}_{\text{LARGE}}$, inference was performed using the span classification method employed in LUKE (Yamada et al., 2020), which classifies whether a span from one word to another is a named entity.

D Pseudo-Incorrect Labeling Method

In the experiments in §5.1, Some of the original labels were replaced with pseudo-incorrect labels. This section describes the replacement method. We replaced the labels so that the replacement would not conflict with the IOB2 format and the number

Method	Subwords	the number of subwords
Default Tokenization	Japan then laid siege to the Syrian penalty area for most of the game but rarely breached the Syrian defence .	20
Random	Japan then la id s i e g e to t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	29
	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	33
	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	30
Cos-Sim	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	43
	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	36
	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	42
K-means	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	31
	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	36
	J a p a n t h e n l a i d s i e g e t o t h e S y r i a n p e n a l t y a r e a f o r m o s t o f t h e g a m e b u t r a r e l y b r e a c h e d t h e S y r i a n d e f e n c e .	34

Table 6: Examples of subword sequences by each selection method. Subword breaks are shown in the blanks.

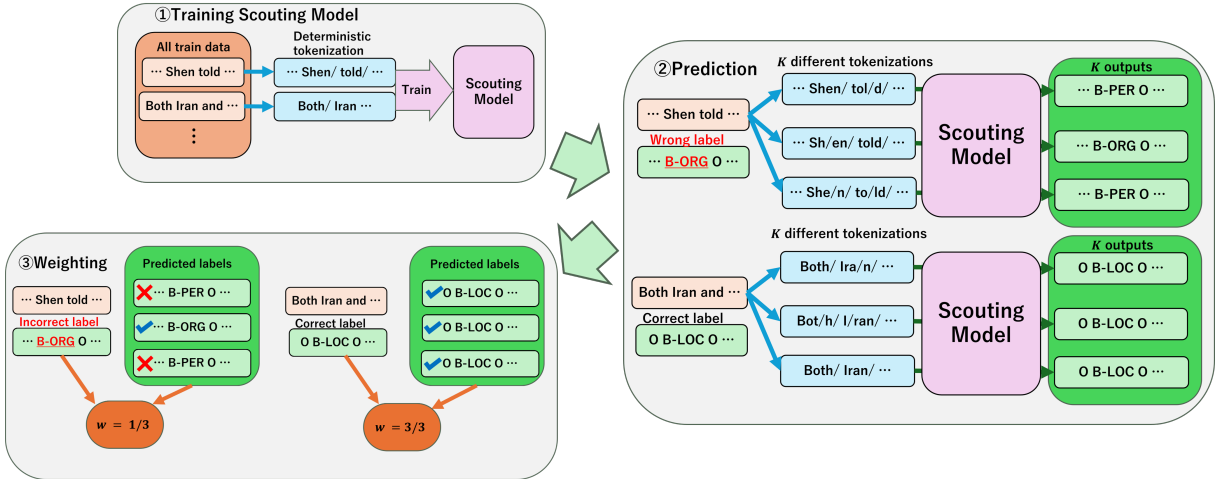


Figure 2: The overview of SubRegWeigh in the case of NER.

of entities would not increase too much from the original data. Specifically, in the case of replacing $a\%$ of all labels, we used the following procedure.

1. Select one label at random.
2. Select one label at random from B-xxx or O that is different from the selected label and replace that label with the selected label.
3. Apply the following operations, depending on the labels to change and to be changed.
 - (a) If changing from O to B-xxx and the label behind the selected label is B-xxx, change it to I-xxx.
 - (b) If changing from B-xxx or I-xxx to O and the label behind the selected label is

I-xxx, change it to B-xxx.

- (c) If changing from B-xxx or I-xxx to B-yyy, change the subsequent I-xxx to I-yyy.

4. Repeat until the number of changed labels reaches $a\%$ of all labels.

E Effect of Subword Regularization

We select a few tokenization candidates from a large number of candidates to use a wide range of different candidates from the original tokenization for the annotation weighing. Instead of using a large number of candidates, we can also obtain various candidates by changing the BPE-Dropout hyperparameter p . Therefore, we investigated the

Model	RoBERTa		LUKE
Task	NER	Text Classification	NER
Epoch	5(20)*	5(20)*	5
Learning rate	1e-5	5e-5	1e-5
Batch size	32	32	16
Weight decay	0.01	0.01	0.01
Params	354M (LARGE)	354M (LARGE)	560M (LARGE)
GPU	V100*1	A100*1	A100*1

Table 7: Model hyperparameters. *In RoBERTa, the scouting model was trained in 5 epochs and the final model was trained in 20 epochs. Other hyper-parameters are the same in the scouting and final model.

p	N	Method	Time	CoNLL test	CoNLL CW
0.2	10	Random	3:26	93.68	95.28
0.2	100	Cos-Sim	3:45	93.42	95.15
0.2	100	K -means	3:58	93.52	95.20
0.1	10	Random	3:26	93.51	95.24
0.1	500	Cos-Sim	4:51	93.44	95.17
0.1	500	K -means	5:21	93.81	95.45

Table 8: Difference in the speed and performance against the hyperparameter of subword regularization p .

difference in F_1 score of the final model against p in the NER dataset. In this examination, we used $p = 0.2$ as the hyperparameter of BPE-Dropout, which tends to generate more different tokenization candidates than $p = 0.1$. the number of tokenization candidates was limited to $N = 100$ to explore the possibility of more efficient annotation weighing by increasing the p of BPE-Dropout to generate diverse subword sequences. For Random, the number of generated subword sequences N always equals the number of subword sequences used for inference K . Since the experiment was conducted with $K = 10$, we set to $N = 10$ in the experiment for Random.

The results are shown in Table 8. One can see that the larger p improves the performance with Random, while the weighing time did not change. This suggests that Random with $p = 0.1$ did not obtain a sufficient variety of tokenization candidates. For non-Random selection methods, the time for weighing annotation errors was reduced to almost equivalent to the ones by Random. However, F_1 score decreased. The non-random selection methods use TF-IDF when comparing subword sequences. Therefore, when selecting from high p and small N , the number of subwords that appear only in a single sequence will increase, and the selection will be attracted by such subwords. In

w_{min}	CoNLL test	CoNLL CW
1 (Vanilla)	93.54	95.27
0.7	93.58	95.23
0.3	93.81	95.45
0	93.25	95.08

Table 9: Difference in the performance against w_{min}

addition, subwords that appear only in a single sequence are more finely segmented subwords, i.e., almost character-level subwords, so that a model trained with deterministic subwords cannot make correct inferences. This is likely the reason for the low f_1 scores of the non-Random selection methods with $p = 0.2$. From these results, it is evident that using a large N with a small p is appropriate for balancing both speed and performance.

F Effect of Minimum Weight

In the experiment shown in §4, the minimum weight w_{min} for SubRegWeigh was set to $1/3$. Herein, we investigate the effect on the F_1 scores of the CoNLL-2003 test data and CoNLL-CW when w_{min} is changed, using RoBERTa LARGE. We use $w_{min} = 1/3, 2/3$, as well as $w_{min} = 0$, where no minimum weight is set. For comparison, Vanilla is recorded with $w_{min} = 1$, which is equivalent to not performing any weight correction, as all data are given 1 as the weight.

The results are shown in Table 9. When $w_{min} = 0$, the F_1 scores decrease compared to other settings. This is likely because the weight of some data becomes 0, reducing the amount of data used for training, and therefore, the model is not sufficiently trained. This is similar to the decrease in CrossWeigh (Wang et al., 2019) when data with low weights were not used instead of being weighted. When $w_{min} = 2/3$, the F_1 scores also decrease compared to $w_{min} = 1/3$. This is likely because the dataset is weighed between $2/3$ and 1, making the dataset almost identical to the vanilla baseline. This result indicates that w_{min} should be set to a low but non-zero value.

G Additional Datasets

We experimented with additional datasets, WNUT2017 (Derczynski et al., 2017) and MRPC (Dolan and Brockett, 2005). We use the train and test split of WNUT2017 and the train and develop split of MRPC. The basic experiment

	WNUT2017	MRPC
	F_1	ACC
SoTA	60.45^{†††}	-
Vanilla	60.04 \pm 0.31	90.43\pm0.15
CrossWeigh	60.19 \pm 0.43	90.35 \pm 0.23
SubRegWeigh		
Random	60.15 \pm 0.30	90.16 \pm 0.23
Cos-Sim	60.05 \pm 0.46	89.69 \pm 0.36
K-means	60.29 \pm 0.41	86.82 \pm 0.45

Table 10: The results of the additional datasets. ^{†††} is SoTA score from Wang et al. (2021c).

setup is the same as for §4.3 and Appendix C, but in WNUT2017, we use only RoBERTa and the URLs in the text are replaced with <URL> tags.

The results are shown in Table 10. In WNUT2017, the proposed method improves for baselines as in the CoNLL2003 experiment. However, in MRPC, the proposed method has worse accuracy than baseline and CrossWeigh. MRPC is a task to classify whether two sentences have the same meaning. the same words are often used in each of these two sentences. If subword regularization splits these words into different subwords, the scouting model cannot perform inference successfully. This is the reason for the accuracy deterioration in MRPC.