

Bit-level BPE: Below the byte boundary

Anonymous ACL submission

Abstract

Byte-level fallbacks for subword tokenization have become a common practice in large language models. In particular, it has been demonstrated to be incredibly effective as a pragmatic solution for preventing OOV, especially in the context of larger models. However, breaking a character down to individual bytes significantly increases the sequence length for long-tail tokens in languages such as Chinese, Japanese, and Korean (CJK) and other character-diverse contexts such as emoji. The increased sequence length results in longer computation during both training and inference. In this work, we propose a simple compression technique that reduces the sequence length losslessly.

1 Introduction

Byte-pair Encoding (BPE) (Sennrich et al., 2016) is a method that allows models to have a robust vocabulary that is capable of representing rare words that have not been seen during training. Variants of this method have been used extensively in many modern natural language processing (NLP) systems, as they allow the representation of a large vocabulary through the concatenation of smaller units, known as subwords, which in turn allows setting an upper boundary on the logits needed for a given model. Byte-level BPE is an extension of BPE to mitigate out-of-vocabulary (OOV). Instead of falling back to OOV when a token cannot be represented through its subwords, it instead encodes the missing token into a sequence of (usually Unicode, in particular - UTF-8) bytes. Due to the introduction of this method, OOV has been largely eliminated in large, foundational models.

However, there is no such thing as a free lunch. Firstly, the model must also learn the intricacies of generating valid Unicode output for any byte-level token, on top of the main linguistic learning (e.g., language modeling) task. For robust generation, training requires enough samples for the

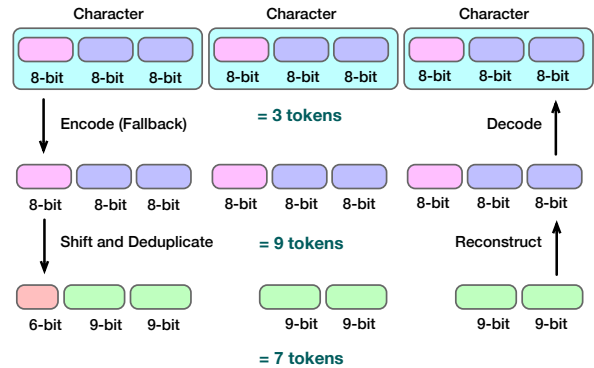


Figure 1: High-level overview of byte-level text representation. In this work, we propose a deduplication method that reduces sequence length at byte level.

model to learn byte sequencing, increasing training costs. Secondly, representing a character in bytes increases the length by up to four times the input. This limits the utility outside of a large model setup.

In this work, we investigate the limitations of byte-level fallbacks, particularly the inefficient nature of the byte-level representation. We propose a tokenizer-agnostic method for reducing redundant information in UTF-8 byte-level fallbacks, which results in shorter sequence lengths - eventually saving compute time. Through experiments across three character-diverse languages (Chinese, Japanese, and Korean), we measure the amount of sequence length reduction and assess the pros and cons of our method. To better understand the experiment results and trade-offs, a method to compute the perceived throughput is proposed for comparison.

2 Background

2.1 Byte-level BPE

Byte-level BPE is an extension of BPE that addresses limitations in dealing with unseen subwords. By adding the lowest-level building block for Unicode characters into the vocabulary, it guar-

antees that there will be no OOV. While it is unclear where byte-level BPE was initially proposed, mainstream usage began as Radford et al. (2019) and Wang et al. (2019) proposed it in the context of language modeling and machine translation. The method increases robustness against OOV, which is a common problem in CJK languages.

This problem is due to the sheer diversity of characters needed to represent CJK languages. The CJK Unified Ideographs block defines a total of 97,680 code points, and the Hangul syllables block defines a total of 11,172 code points¹. Any form of character-level bigram merge, therefore, will result in a combinatorial explosion, which in turn makes supporting any of these languages computationally costly. This complexity, combined with the relatively smaller amount of textual data available in the wild compared to other languages, resulted in many of the CJK characters being treated as OOV. By using byte-level fallback instead of naively losing information (e.g., OOV), modern models have effectively mitigated this problem.

2.2 Related Work

Our work is related to the history of character-level and byte-level NLP methods (Zhang and LeCun, 2017; El Boukkouri et al., 2020; Shaham and Levy, 2021; Xue et al., 2022). Although these methods using fine-grained tokenization alleviate the problem of unknown words, they do so at the cost of performance on longer inputs (Libovický and Fraser, 2020; Gowda and May, 2020; Goldman et al., 2024). Byte-level NLP incurs even more severe performance penalties with longer inputs (Mielke et al., 2021; Sreedhar et al., 2023). Rust et al. (2021) analyzes the effect of over-segmented input, particularly its adversarial effects on model performance.

In the recent LLM era, the fairness of unequal LLM usage costs among different languages is discussed (Ahia et al., 2023; Petrov et al., 2024). Our proposed method can shorten the input length while preserving byte-level fallback, alleviating disparities in training and inference costs among languages.

2.3 Byte-level Limitations

The clear benefit of using byte-level is that it guarantees there will be no OOV in the final trained model. This benefit, however, comes at the cost of

¹Statistics as of Unicode 15.1.

three major caveats.

The first caveat is that one needs to have enough training to ensure that the model does not generate invalid Unicode sequences for a sequence containing rare tokens, as without having seen enough samples, the model will inherently generate invalid Unicode sequences. This results in an implicit requirement to have a larger training dataset to ensure better generalization. This tends to indirectly affect the model size, as the size of the dataset and model should ideally be proportional (Hoffmann et al., 2022) - increasing the training cost of the model.

The second is the increased sequence length, which is also inversely proportional to the size of the model’s vocabulary². On average, for a CJK character to be represented with UTF-8 bytes, the sequence length grows by three times compared to a character-level representation. The smaller the vocabulary is, the higher the probability of OOV vocabulary is, therefore the sequence length increases. There are tradeoffs to be made here, as a larger vocabulary increases the cost of logit computation and, as a result, requires more computation power for training and inference - but also can introduce undertrained tokens (Land and Bartolo, 2024).

The last one is a bit more subtle and relates to the tokenizer quality against the byte-level token-only distribution. Tokenizers are often trained with a small sample of the final training corpus, which results in a suboptimal distribution when observed from a whole-corpus perspective. As an unintended side effect, it also creates a local distribution specific to the byte portion, which has distributional characteristics of the underlying character reflect³. Zouhar et al. (2023) proposes Rényi efficiency as a measure of tokenization quality. This is computed through Rényi entropy $H_\alpha(W_v)$ of the random variable W_v distributed according to p_v for a given vocabulary V , defined as:

$$H_\alpha(W_v) = \lim_{\alpha' \rightarrow \alpha} \frac{1}{1 - \alpha'} \log \left(\sum_{w \in V} p_v(w)^{\alpha'} \right) \quad (1)$$

and the associated Rényi efficiency $E_\alpha(W_v)$ is a scaled with respect to the vocabulary size $|V|$:

$$E_\alpha(W_v) \approx \frac{H_\alpha(W_v)}{\log |V|} \quad (2)$$

²It is somewhat proportional to the number of languages one needs to support. The more languages supported in a single vocabulary, the more likely there will be a strong dependency on byte-level fallbacks.

³In the case of most models, this is UTF-8.




Figure 2: Byte-level frequency distribution on the Japanese-Korean subtitle translation task training set, ranked by most frequent byte.

In this context, a sharp probability mass distribution of token frequency results in lower Rényi entropy, and intuitively, the more uniform the distribution is, the more entropy increases. Efficiency is a normalized form of entropy, accounting for the vocabulary size $|V|$. In our experiment and tokenizer setup, we can assume that the byte-specific distribution has low entropy, as the frequency disparity between the most frequently-observed bytes and the rest of the vocabulary⁴ is high, resulting in suboptimal tokenizer efficiency when observed at the byte-level portion of the token distribution. To minimize misunderstandings, we will refer to Rényi efficiency as **entropy** throughout this paper.

3 Proposed Method

We propose an efficient tokenization method to shorten the input sequence, focusing on the repetition (§3.1) and duplication (§3.2) in UTF-8. This section first provides the observation nature of the general representation of texts in UTF-8 and then explains the proposed method, which provides the efficient representation (§3.3). Figure 1 overviews the proposed method.

3.1 Duplication in UTF-8

The high-frequency token phenomena observed in Figure 2 is a bias inherited from the underlying character encoding, in this case, UTF-8. For the examples shown in Table 1, the UTF-8 encoded byte sequence from Chinese, Korean, and Japanese exhibits a duplication problem, indicated in bold. By simply observing this small sample, it is clear that byte-level tokens in the range of E4-ED have a much higher frequency than others.

⁴In this case, all 255 possible bytes.

zh-CN	召	喚	众
	E4 BC 97	E5 94 A4	E4 BC 97
ja-JP	検	認	裁
	E6 A4 9C	E8 AA 8D	E8 A3 81
ko-KR	철	저	히
	EC B2 A0	EC A0 80	ED 9E 88

Table 1: Example text from Chinese, Japanese, and Korean, accompanied with their byte representations. Each character needs three tokens to represent, and we also observe that there are common prefixes.

E4	E5	E6	E7
111001 00	111001 01	111001 10	111001 11
E8	E9	EA	EB
111010 00	111010 01	111010 10	111010 11
EC	ED	EE	EF
111011 00	111011 01	111011 10	111011 11

Table 2: UTF-8 CJK block common prefixes as seen in bits. E4-E9 are CJK ideographs, while EA-ED is Korean.

As UTF-8 must guarantee robustness to encode every possible character in a deterministic form, duplication is encoded into the scheme so it can be decoded without dependency on the context. This is an important property for reliable transmission and storage, but it has several undesirable properties in the context of text generation from raw bytes. Out of the three undesirable properties discussed in 2.3, the two we intend to address directly in our work are the duplication of information and learning complexity. The length problem is as demonstrated in Table 1, and the learning complexity can be attributed to the duplication.

3.2 Duplication at Bit-level

Looking at the bit level, the duplication of byte representation can be generalized as the duplication of bit sequences. We observe that the underlying frequency distribution of the byte-level token range shows overly frequent tokens, supporting our duplication claim. Not only are a third of the bytes redundant information, as shown in Table 1, but we hypothesize that the redundancy also increases the complexity of learning, as it requires learning the linguistic structure while also learning to generate valid UTF-8 sequences, including emitting the redundant tokens⁵. We also observe a commonality

⁵A somewhat simple analogy of this is learning how to count while also learning how to play Fizz buzz.

between the most frequent tokens at the bit-level observed in Table 2.

Out of the three text samples in Table 2, we can observe that aside from the Japanese (ja) sample, the first six bits of the prefix - 0xEC-0xEF for Korean (ko), and 0xE4-0xE7 for Chinese (zh) are shared. This suggests that there is a possibility that this could be considered redundant information that can be de-duplicated.

However, to make this de-duplication possible, we must break a common assumption - that a byte representation, as expressed by the model’s tokenizer, is constrained to be eight bits. Having byte-level tokens represented as an exact byte counterpart has an advantage - as long as the model can generate valid byte sequences, it can be decoded without any extra effort. However, this only holds if the model is trained with enough data to learn the linguistic constructs of the corpus and how to generate valid UTF-8. Unless training a large model, this is very challenging, and even in a large model context, it increases the cost of training.

3.3 Breaking the 8-bit Byte Boundary

Overview

Our method breaks this common assumption, eight bits for a byte, by treating a sequence of bytes as a sequence of bits. Instead, ignoring the eight-bit boundary and having a flexible boundary of bits allows us to optimize the sequence - in particular, by removing duplicate information. With our method applied, a "byte" in a tokenizer can hold any amount of bits instead of being uniformly eight (see Figure 1). This is possible as the model sees the individual bytes as logits, which does not require them to be aligned to eight bits. This optimization would not be possible outside of this context without wasting space.

With bit boundary constraints lifted and the common bits identified, we can now treat the common bits as a shared bit prefix of the upcoming tokens. For example, we can treat 0b111001 (0x39) to indicate that all following tokens have the same prefix unless the prefix changes or a subword is emitted. The two residual bits (e.g., 01 in the case of E5) will need to be carried over by the following bytes. As the vocabulary allocated for bytes is 2^8 , carrying the two bits over on the next byte will increase the vocabulary budget significantly - as this would mean adding 768 more "byte" representation tokens ($2^{10} - 2^8$). For this reason, we redistribute

the bits to have two 9-bit integers for our new extended "byte" representation. This requires adding 256 more tokens.

We demonstrate this proposed method with the Chinese input shown in Table 1 (召唤众):

E4 BC 97 E5 94 A4 E4 BC 97

The process is divided into two parts: encoding and decoding of the text.

Encoding

To isolate the common bits, we choose the first "byte" to be six bits, which results in two residual bits - for example with the character E4 BC 97, the bit boundaries move as follows:

```
11100100 10111100 10010111
111001 00 10111100 10010111
111001 001011110 010010111
```

After the bit re-distribution, the first six bits are now $p1=0b111001$, a special token representing a 6-bit prefix⁶. This can be implemented through a simple set of bitwise operations; given three bytes b_1 , b_2 , and b_3 representing a single character, the prefix \hat{b}_1 , and 9-bit tokens \hat{b}_2 , and \hat{b}_3 , are computed by⁷:

$$\hat{b}_1 = (b_1 \wedge 127) \gg 2 \quad (3)$$

$$\hat{b}_2 = (((b_1 \wedge 3) \ll 7) \vee (((b_2 \wedge 254) \gg 1) \quad (4)$$

$$\hat{b}_3 = ((b_2 \wedge 1) \ll 8) \vee b_3 \quad (5)$$

By shifting the existing byte boundaries to be six, nine, and nine bits, respectively, we can now represent the same sequence as:

p1 5E 97 p1 CA A4 p1 5E 97

With naive incremental encoding and a deterministic trailing sequence length for each character⁸, instances of p1 can now be de-duplicated, as the entire sequence is behind the first shared 6-bit prefix. Note that a special token is used instead of the byte representation (0x39) to disambiguate between naturally occurring bytes and incremental decoder triggers. The de-duplicated sequence is:

p1 5E 97 CA A4 5E 97

This is a 22.22% reduction from the original input. Whenever there is a transition in the 6-bit

⁶ $p1=0x39$, which is duplicated across the sequence. In practice, all prefix tokens can usurp existing byte tokens, as they are unreachable. For example, 0x39 is the character "9".

⁷ \ll and \gg are bitwise left and right shift, respectively.

⁸A prefix is followed by two trailing bytes.

		Length ↓			Entropy ↑	
		Byte	Ours	Diff	Byte	Ours
en-zh	T	131M	127M	3.13%	0.764	0.634
	B	64M	60M	6.41%	0.586	0.435
en-ja	T	176M	174M	0.83%	0.818	0.763
	B	41M	40M	3.56%	0.580	0.417
ja-ko	T	113M	111M	2.21%	0.498	0.485
	B	59M	56M	4.25%	0.485	0.446

Table 3: Sequence length reductions across the training sets, rounded to the nearest million tokens. T indicates the entire corpus, and B indicates the byte-level portion.

prefix, the new 6-bit prefix p_n is emitted. For example, if we append a prefix transitioning Japanese character (認) E8 AA 8D, the resulting sequence (召喚众認) with our method is encoded as:

p1 5E 97 CA A4 5E 97 **p2** 55 8D

With one prefix switch, the gains decrease - 16.66% shorter than the original input.

Decoding

As the specific Unicode blocks our method targets have a deterministic length of three bytes per character, we can invert the incremental encoding by emitting the current prefix token after every 9-bit bi-gram, which results in the following generated sequence:

p1 5E 97 **p1** CA A4 **p1** 5E 97 **p2** 55 8D

This sequence can be re-aligned to a boundary of eight bits each, which results in a decodable UTF-8 byte sequence, through the following computation:

$$b_1 = (\hat{b}_1 \ll 2) \vee (\hat{b}_2 \gg 7) \quad (6)$$

$$b_2 = ((\hat{b}_2 \wedge 127) \ll 1) \vee (\hat{b}_3 \gg 8) \quad (7)$$

$$b_3 = \hat{b}_3 \wedge 255 \quad (8)$$

This results in the reconstructed sequence:

E4 BB BD **E5** 81 87 **E7** AE 80 **E8** 94 B5

In our work, the bit boundaries were set to be optimal for CJK scripts - the boundaries can be set differently for different Unicode blocks.

4 Experiments

To observe the effects of our method, we experimented across three target languages (Chinese, Japanese, and Korean) on a translation task.

4.1 Datasets

We used three datasets for this experiment, referenced by their target language for simplicity throughout the paper: English-Chinese (Chinese),

English-Japanese (Japanese), and Japanese-Korean (Korean). The Chinese and Japanese datasets used are from WMT20 (Barrault et al., 2020), while the Korean dataset is from AI Hub⁹. For the WMT20 shared task, we used a custom split of the Wiki-Matrix dataset from the news translation task for Chinese (2.3M) and Japanese (3.5M). For Korean, we used a subtitle translation dataset (3M)¹⁰. Each dataset had a 5K test set held out for evaluation and performance benchmarking and 40K for validation. The remainder was used for training.

4.2 Model Settings

Each experiment trains byte-level fallback on a machine translation (MT) model trained from scratch. We use a pre-trained Llama2 tokenizer, which triggers a significant amount of byte-level fallbacks in CJK languages, as can be seen in Table 3 (see §4.3). We compare a Llama2 tokenizer for baseline byte-level BPE (**Byte**) to an augmented Llama2 tokenizer with our method (**Ours**). We consider this popular tokenizer to be a valid baseline to examine our method, which alleviates the problem in the byte-level representation. All experiments use identical vocabulary expansion when applying our method - our method requires 256 more byte-level tokens (0x100-0x1FF), along with three prefix tokens (p_1 , p_2 , p_3).

The comparison intends to approximate the effects of our method in a from-scratch pre-training setup under a limited compute budget. Specifically, this experiment aims to demonstrate the challenges of learning byte-level generation on a small model and observe the effects of our method when generating byte sequences compared to a baseline byte-level BPE tokenizer. Here, we train a 65M parameter vanilla Transformer (Vaswani et al., 2017). With the different tokenizations, we trained the model for a fixed number of comparable epochs¹¹.

Only translation into the CJK languages is evaluated, as the generation of CJK text has a critical dependency on first being able to generate a valid UTF-8 sequence. We expect the trained model to underperform, as we hypothesize that byte-level learning in a small model context is challenging.

The model quality was evaluated after reconstructing the byte sequence and performing another

⁹<https://aihub.or.kr>

¹⁰An English-Korean dataset of comparable size was not available for free use as of the time of writing.

¹¹With an exception where we stop the model training if it stalled for 10 epochs in a row.




Figure 3: Byte-level distribution shift, with and without our method on the Japanese-Korean subtitle translation task. We can observe that while the probability mass is distributed in our proposed method, there is a sharp increase in the frequency of the most frequent tokens ($0 \times 39 - 0 \times 3C$). This results in lower entropy.

	Decode error ↓		Empty ↓	
	Byte	Ours	Byte	Ours
en-zh	33	14	3156	2545
en-ja	136	0	218	87
ja-ko	1522	121	7	3

Table 4: Number of invalid output from the 5000 test samples.

detokenization step. With the resulting output, we used sacreBLEU (Post, 2018) under different configurations to compute the translation scores. Different configurations were used for each language, and a unified configuration was used to compute byte-level BLEU and chrF. The sacreBLEU signatures are disclosed in our Appendix. For throughput benchmarks in Table 6, we measured the wall time for inference against 5000 samples, with a batch size of 1000¹².

We use a combination of Marian (Junczys-Dowmunt et al., 2018), and Huggingface Transformers (Wolf et al., 2020) for these experiments.

4.3 Results

Sequence length and Entropy

Table 3 reports the sequence length of the texts tokenized with the original Llama2 tokenizer (Byte) and the one with the proposed method (Ours). This table demonstrates that the Llama2 tokenizer yields a lot of byte-level tokens for each dataset (e.g., 64M

¹²This is not an entirely realistic setup, as inference time batches will likely not be computed in batches of 1000.

out of 131M tokens are byte-level in the “en-zh” dataset tokenized with the baseline). This result also shows that the proposed method shortens the length of sequence across all CJK datasets with the de-duplication technique, which results in increased computation efficiency. The magnitude of sequence reduction was higher in Chinese and Korean. The reason for this can be explained through the proportion of byte fallback, Japanese only has 23.4% byte fallback, while Chinese has 48.8% and Korean as 52%. This sequence length comes at a cost - the immediate side effect being an increase in parameter count.

Table 3 shows changes in entropy as computed by Renyi efficiency, defined in equation (2). As our method’s application decreases tokenization’s entropy, we can assume that there is a model quality degradation - as Zouhar et al. (2023) found entropy correlates to quality. In particular, as observed in Table 3 and Figure 3, a trade-off between entropy and sequence length needs to be made. In our work, we prioritize sequence length over entropy.

Reduction of Decoding Error

The number of sequences that were unable to be decoded see a significant decrease, as can be observed in Table 4. We also see a significant decrease in empty sequences, although, for Chinese, more than half of the generated output was empty¹³. This required us to train the Chinese model for 100

¹³The initial model (200 epochs) was much worse, with 95% of the output of the 5K test samples given translated to blank text on both byte and ours.

	Size	BLEU \uparrow		chrF \uparrow		TER \downarrow		Byte BLEU \uparrow		Byte chrF \uparrow	
		Byte	Ours	Byte	Ours	Byte	Ours	Byte	Ours	Byte	Ours
en-zh	2.23M	0.5	1.7	1.9	3.2	127.3	100.4	3	7.3	3	5.5
en-ja	3.47M	1.6	3.2	6.5	8.5	467.2	100.3	15.2	20	19.4	23.5
ja-ko	2.99M	19	24.6	35.8	35.5	114.4	100.2	49	53.3	54.7	53.5

Table 5: Translation performance results across the three language pairs from 5000 test samples. Byte BLEU and Byte chrF operates were computed at the byte level after reconstructing the output sequence. The table heading \uparrow indicates metrics where higher is better, while \downarrow it indicates metrics where lower is better.

	en-zh		en-ja		ja-ko	
	Byte	Ours	Byte	Ours	Byte	Ours
Tokens Out	188,361	194,891	578,446	465,401	168,233	302,867
AvgTok Out	33.64	60.57	115.69	93.08	37.67	38.98
Total time (s)	72.41	201.65	529.05	271.60	41.10	41.59
Tokens per Second (TPS)	464.69	300.39	183.91	342.71	916.43	937.21
Tokens in test reference	291,857	282,423	253,540	251,261	189,800	185,628
Relative Gain	1	1.0334	1	1.00091	1	1.0225
Perceived TPS	464.70	310.43	183.91	345.81	916.43	958.28

Table 6: Perceived TPS of a model. Tokens out are the total tokens output during the test, and AvgTok is the total divided by the test set size (5K). Tokens per Second were computed with the total tokens divided by the mean runtime across 5 runs. Relative gain is the perceived TPS improvement by reduced sequence length.

more epochs, significantly reducing the amount of blank output. Other languages tended to have many invalid sequences; Korean had a higher tendency to produce invalid byte sequences over blank sequences, while Japanese suffered the least. This is likely because Japanese has the best in-vocabulary coverage out of the three languages, requiring less dependency on byte-level fallbacks, as seen in Table 3. These results suggest that byte-level generation is challenging for smaller model setups.

Task Performance

Table 5 shows the translation performance results. The first thing that is very clear is that the BLEU scores are quite distant from public baselines based on the trained models. In particular, Chinese and Japanese are unlikely to be meaningful candidates for quantitative comparison. For this reason, we will focus on the Korean model, where we can perform a better analysis. For the Korean task, we see modest gains between our proposed method and the byte baseline in every aspect. However, this comes at the cost of longer computing times. We see gains for Chinese and Japanese, but both models are likely to be severely undertrained.

We hypothesize that using the Transformer architecture with the default sequence length, long inputs, and byte-level tokenization may have con-

tributed to the disappointing BLEU scores. The primary difference between the cases where we observe positive signals compared to the undertrained model cases is the average length per sample - as subtitles tend to have much shorter sentences.

5 Factoring Tokenization into TPS

Table 6 reports the statistics about the inference time in each dataset. Initially, we naively assumed that sequence length is linearly correlated to the wall clock time needed by the accelerator to compute the sequence. This did not turn out to be true, and the extra token cost especially needs to be factored in, as it introduces extra computation with an increase in parameters. There are also extra costs when computing the logits and attention. This means one needs to consider the gains that come from reducing the sequence length with our method and whether the extra compute cost compared to the gains results in a net loss. This was one area where there was not much existing literature, which required us to invent an approximation method of the perceived TPS.

Perceived TPS, simply put, is the time it takes to convey the same amount of text to a recipient with different tokenization and model configurations. The coefficients needed to approximate the perceived TPS are in Table 6. The most common

489 measure of throughput used is tokens per second
490 (TPS), which is a reliable measure of the model’s
491 output speed. However, it does not consider the
492 inefficiencies stemming from poor tokenization. A
493 byte-heavy output can require around twice the
494 amount of tokens generated to transmit the same
495 amount of textual information compared to that of
496 a model that does not use bytes. Our method re-
497 duces the sequence length, reducing the amount of
498 tokens needed to convey the same information.

499 To compute perceived TPS, we introduce the
500 concept of *relative gain*, a simple TPS multiplier
501 that factors in the expected differences between the
502 two tokenizers of models with identical TPS. This
503 is computed as the relative sequence length of the
504 same text between two tokenizers. The assumption
505 made here is that given the same text, a tokenizer
506 that is capable of encoding it with less tokens has
507 higher expressive power, therefore when used in
508 conjunction with TPS results in a better throughput
509 score of the model. It is computed by $\frac{|T_c|}{|T_e|}$, where T_c
510 and T_e are tokenized sequences using the control
511 and experimental tokenizers, respectively.

512 For our results, we used the test set for each lan-
513 guage to compute the gain for our case, assuming
514 optimal output. This was a conscious decision, as
515 the relative gain computed by a poorly performing
516 model output tended to overamplify the gains¹⁴.

517 For example, the gain of our method on the Ko-
518 rean translation task is 1.0334, which suggests that
519 there is a 3.34% advantage in expressive power
520 with the same token count compared to the base-
521 line. This can be multiplied by TPS better to ap-
522 proximate the model and tokenizer combination’s
523 perceived TPS. In the last row of Table 6, perceived
524 TPS can be compared to conventional TPS.

525 6 Conclusion

526 Byte-level representation is a simple yet effective
527 method for ensuring full vocabulary coverage for
528 arbitrary Unicode input. However, it comes at the
529 cost of longer sequences, which increase computa-
530 tion time in training and inference. Additionally, it
531 has a high risk of generating invalid sequences, es-
532 pecially with smaller models and training datasets.
533 Our work shows that the effects of these undesir-
534 able properties can be partially mitigated through
535 a minor re-encoding of input and output data. Our
536 proposed method is validated by demonstrating its

¹⁴This was mostly caused by repeated output in under-
trained model output.

537 efficacy in reducing sequence length and the num-
538 ber of failed decode operations. The positive results
539 demonstrate that raw UTF-8 bytes are a subopti-
540 mal representation and suggest further investigation
541 into an alternative method for text representation,
542 especially at the sub-byte scale.

543 While evaluating this work, we discovered that
544 TPS, as a measurement of model throughput, is
545 not representative of throughput when seen from
546 a completed text output perspective. We hope that
547 in future model throughput discussions, factoring
548 in tokenization through relative sequence length
549 between models will be used in conjunction with
550 TPS. As it is presented today, the findings in this
551 work have a theoretical trade-off concerning tok-
552 enizer entropy and sequence length. Additionally,
553 the work is still constrained to the underlying UTF-
554 8 framework. This is still far from optimally han-
555 dling long-tail tokens. We expect our findings to be
556 useful as an incremental step toward further chal-
557 lenging the status quo of naively using UTF-8 for
558 long-tail token coverage.

559 7 Future Work

560 Byte-level and particularly alternative encoding
561 methods are underexplored as of today, and many
562 promising avenues of investigation remain.

563 In this work, the bit spans are fixed for imple-
564 mentation simplicity reasons. This is an acceptable
565 first step to validate this method for the context of
566 our experimental setup, scoped to CJK languages.
567 However, this can be further expanded as an opti-
568 mization problem - minimizing sequence length,
569 minimizing the amount of extra tokens needed, and
570 maximizing entropy. Our work does not address
571 the problem of finding optimal bit boundaries.

572 Our method decreases tokenization entropy,
573 likely degrading performance. This might be fix-
574 able by decreasing the frequency of the most fre-
575 quent tokens. For example, one could omit the pre-
576 fix if the previous subword shares the same prefix.

577 Additionally, the increase in model size caused
578 by the extra vocabulary can be further optimized
579 by utilizing unreachable and, therefore, untrained
580 byte-level tokens. We considered this a premature
581 optimization and did not explore this direction.

582 Finally, the utility of perceived TPS was only
583 investigated for the scope of our work. We expect
584 this to be more useful when applied to configura-
585 tions with a larger delta, such as when comparing
586 Llama2 (32K) and Llama3 (128K), for example.

587 Limitations

588 Our work is an early investigation in a previously
589 underexplored area of tokenization. As this is the
590 first step in a new direction for byte-level tokeniza-
591 tion methods, we focused mostly on the paradigm
592 shift of disregarding byte boundaries. As noted in
593 Section 7, the investigation was simplified by re-
594 stricting the scope to CJK languages, allowing us
595 to have a fixed boundary. This particular split may
596 not be optimal for other scripts.

597 Dataset limitations and compute costs were huge
598 factors that limited the scope of our investigation.
599 CJK languages are comparatively higher-resource,
600 compared to scripts in some other larger blocks
601 - such as Tangut (6,904 characters) or Yi (1,220
602 characters) where resources are scarce in terms
603 of both datasets and raw web corpora. Emojis are
604 another area we excluded from the investigation, as
605 they are likely not as frequent as text and, therefore,
606 unlikely to benefit as much from deduplication.

607 Ethical Statement

608 Our method itself does not introduce any new eth-
609 ical risks when used in conjunction with existing
610 methods. The method is data and task-agnostic
611 and is a method of lossless sequence compression.
612 There are likely minor environmental benefits that
613 come from the reduced computing costs.

614 The models from our experiments carry the same
615 ethical risks and bias as the underlying dataset used
616 to train them. To minimize the risk, we resorted to
617 well-known or openly available datasets whenever
618 possible and, therefore, can state that we are intro-
619 ducing no novel vectors of societal harm through
620 this work.

621 References

622 Orevaoghene Ahia, Sachin Kumar, Hila Gonen, Jungo
623 Kasai, David Mortensen, Noah Smith, and Yulia
624 Tsvetkov. 2023. [Do all languages cost the same?
625 tokenization in the era of commercial language mod-
626 els](#). In *Proceedings of the 2023 Conference on Em-
627 pirical Methods in Natural Language Processing*,
628 pages 9904–9923, Singapore. Association for Com-
629 putational Linguistics.

630 Loïc Barrault, Magdalena Biesialska, Ondřej Bo-
631 jar, Marta R. Costa-jussà, Christian Federmann,
632 Yvette Graham, Roman Grundkiewicz, Barry Had-
633 dow, Matthias Huck, Eric Joanis, Tom Kocmi,
634 Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof
635 Monz, Makoto Morishita, Masaaki Nagata, Toshi-
636 aki Nakazawa, Santanu Pal, Matt Post, and Marcos

Zampieri. 2020. [Findings of the 2020 conference on
machine translation \(WMT20\)](#). In *Proceedings of
the Fifth Conference on Machine Translation*, pages
1–55, Online. Association for Computational Linguis-
tics. 637
638
639
640
641

Hicham El Boukkouri, Olivier Ferret, Thomas Lavergne,
Hiroshi Noji, Pierre Zweigenbaum, and Jun’ichi Tsu-
jii. 2020. [CharacterBERT: Reconciling ELMo and
BERT for word-level open-vocabulary representa-
tions from characters](#). In *Proceedings of the 28th
International Conference on Computational Linguis-
tics*, pages 6903–6915, Barcelona, Spain (Online).
International Committee on Computational Linguis-
tics. 642
643
644
645
646
647
648
649
650

Omer Goldman, Avi Caciularu, Matan Eyal, Kris Cao,
Idan Szpektor, and Reut Tsarfaty. 2024. [Unpacking
tokenization: Evaluating text compression and its
correlation with model performance](#). *arXiv preprint
arXiv:2403.06265*. 651
652
653
654
655

Thamme Gowda and Jonathan May. 2020. [Finding the
optimal vocabulary size for neural machine transla-
tion](#). In *Findings of the Association for Computa-
tional Linguistics: EMNLP 2020*, pages 3955–3964,
Online. Association for Computational Linguistics. 656
657
658
659
660

Masanori Hirano, Masahiro Suzuki, and Hiroki Sakaji.
2023. [Ilm-japanese-dataset v0: Construction of
Japanese Chat Dataset for Large Language Models
and its Methodology](#). 661
662
663
664

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch,
Elena Buchatskaya, Trevor Cai, Eliza Rutherford,
Diego de Las Casas, Lisa Anne Hendricks, Johannes
Welbl, Aidan Clark, Tom Hennigan, Eric Noland,
Katie Millican, George van den Driessche, Bogdan
Damoc, Aurelia Guy, Simon Osindero, Karen Si-
monyan, Erich Elsen, Jack W. Rae, Oriol Vinyals,
and Laurent Sifre. 2022. [Training compute-optimal
large language models](#). *Preprint*, arXiv:2203.15556. 665
666
667
668
669
670
671
672
673

Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-
Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu
Chen. 2022. [LoRA: Low-rank adaptation of large
language models](#). In *International Conference on
Learning Representations*. 674
675
676
677
678

Marcin Junczys-Dowmunt, Roman Grundkiewicz,
Tomasz Dwojak, Hieu Hoang, Kenneth Heafield,
Tom Neckermann, Frank Seide, Ulrich Germann,
Alham Fikri Aji, Nikolay Bogoychev, André F. T.
Martins, and Alexandra Birch. 2018. [Marian: Fast
neural machine translation in C++](#). In *Proceedings of
ACL 2018, System Demonstrations*, pages 116–121,
Melbourne, Australia. Association for Computational
Linguistics. 679
680
681
682
683
684
685
686
687

Sander Land and Max Bartolo. 2024. [Fishing for
magikarp: Automatically detecting under-trained
tokens in large language models](#). *Preprint*,
arXiv:2405.05417. 688
689
690
691

Jindřich Libovický and Alexander Fraser. 2020. [To-
wards reasonably-sized character-level transformer](#) 692
693

694	NMT by finetuning subword systems . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 2572–2579, Online. Association for Computational Linguistics.	751
695		752
696		753
697		
698		
699	Sabrina J Mielke, Zaid Alyafeai, Elizabeth Salesky, Colin Raffel, Manan Dey, Matthias Gallé, Arun Raja, Chenglei Si, Wilson Y Lee, Benoît Sagot, et al. 2021. Between words and characters: A brief history of open-vocabulary modeling and tokenization in nlp. <i>arXiv preprint arXiv:2112.10508</i> .	754
700		755
701		756
702		757
703		758
704		759
705	Aleksandar Petrov, Emanuele La Malfa, Philip Torr, and Adel Bibi. 2024. Language model tokenizers introduce unfairness between languages. <i>Advances in Neural Information Processing Systems</i> , 36.	760
706		761
707		762
708		763
709	Matt Post. 2018. A call for clarity in reporting BLEU scores . In <i>Proceedings of the Third Conference on Machine Translation: Research Papers</i> , pages 186–191, Brussels, Belgium. Association for Computational Linguistics.	764
710		765
711		766
712		767
713		768
714	Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners .	769
715		770
716		771
717	Phillip Rust, Jonas Pfeiffer, Ivan Vulić, Sebastian Ruder, and Iryna Gurevych. 2021. How good is your tokenizer? on the monolingual performance of multilingual language models . In <i>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)</i> , pages 3118–3135, Online. Association for Computational Linguistics.	772
718		773
719		774
720		775
721		776
722		777
723		778
724		779
725		780
726	Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units . In <i>Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.	781
727		782
728		783
729		784
730		785
731		786
732		787
733	Uri Shaham and Omer Levy. 2021. Neural machine translation without embeddings . In <i>Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies</i> , pages 181–186, Online. Association for Computational Linguistics.	788
734		789
735		790
736		791
737		792
738		793
739	Makesh Narsimhan Sreedhar, Xiangpeng Wan, Yu Cheng, and Junjie Hu. 2023. Local byte fusion for neural machine translation . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 7199–7214, Toronto, Canada. Association for Computational Linguistics.	794
740		795
741		796
742		797
743		798
744		799
745		800
746	Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need . In <i>Advances in Neural Information Processing Systems</i> , volume 30. Curran Associates, Inc.	801
747		802
748		
749		
750		
	Changhan Wang, Kyunghyun Cho, and Jiatao Gu. 2019. Neural machine translation with byte-level subwords . <i>Preprint</i> , arXiv:1909.03341.	751
		752
		753
	Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations</i> , pages 38–45, Online. Association for Computational Linguistics.	754
		755
		756
		757
		758
		759
		760
		761
		762
		763
		764
		765
	Linting Xue, Aditya Barua, Noah Constant, Rami Al-Rfou, Sharan Narang, Mihir Kale, Adam Roberts, and Colin Raffel. 2022. ByT5: Towards a token-free future with pre-trained byte-to-byte models . <i>Transactions of the Association for Computational Linguistics</i> , 10:291–306.	766
		767
		768
		769
		770
		771
	Xiang Zhang and Yann LeCun. 2017. Which encoding is the best for text classification in chinese, english, japanese and korean? <i>arXiv preprint arXiv:1708.02657</i> .	772
		773
		774
		775
	Vilém Zouhar, Clara Meister, Juan Gastaldi, Li Du, Mrinmaya Sachan, and Ryan Cotterell. 2023. Tokenization and the noiseless channel . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 5184–5207, Toronto, Canada. Association for Computational Linguistics.	776
		777
		778
		779
		780
		781
		782
	A Appendix	783
	A.1 Fine-tuning and Undertrained Tokens	784
	As a negative result, we also ran experiments to observe the effects of our method when applied to fine-tuning or continual pre-training of a model. In this experiment, we fine-tuned a Llama2 7B model to observe the effects of our method ¹⁵ . This experiment intends to approximate the effects in a foundational model setting by fine-tuning a model through low-rank adaptation (Hu et al., 2022). We expected this to have characteristics similar to those of pre-training a model from scratch.	785
		786
		787
		788
		789
		790
		791
		792
		793
		794
	Here, we apply our method to the tokenization stage of a pre-trained Llama2 7B model. For newly added tokens, we apply an embedding copy method. This is for better initialization compared to a randomly initialized embedding. A Japanese instruction-tuning dataset with 9.07M training instances (Hirano et al., 2023) was used to fine-tune the pre-trained model.	795
		796
		797
		798
		799
		800
		801
		802

¹⁵This experiment had to be abandoned due to compute budget constraints.

803 The embeddings are copied from the existing
804 byte-level tokens with the smallest hamming distance
805 to the newly added token. For example,
806 $0x1AF$ will be initialized with the same embedding
807 as $0xAF$. The prefix tokens p_1 , p_2 , p_3 get
808 a slightly different treatment, where the embedding
809 is an average across the four prefixes the new
810 prefix token represents. For example, p_1 will be
811 initialized by the mean between the embeddings
812 for $0xE4-0xE7$.

813 This model did not produce any meaningful results
814 after training for 20K steps, which can potentially
815 be attributed to severely undertrained tokens
816 being copied over, causing large losses and, therefore,
817 catastrophic forgetting in the initial steps. We
818 later confirmed this through inspection, and the
819 undertrained tokens can be observed in Figure 4.

820 However, this suggests that these unreachable
821 tokens can be recycled for the newer 9-bit "bytes"
822 ($0x100-0x1FF$) if the parameter budget needs to be
823 optimized.

824 A.2 Environment and Training Setup

825 All of the translation experiments were performed
826 on a shared environment, using two Nvidia H100
827 HBM2 (94GB). The LoRA experiment was done
828 in the same environment, using four of the same
829 compute accelerators. At inference time, all experiments
830 involving wall-clock measurements were
831 done using a dedicated compute node with a single
832 Nvidia A6000 (48GB).

833 Korean translation training was run for 24 hours,
834 Japanese for 72 hours, and Chinese for 96 hours
835 (H100x2). The LoRA experiment (H100x4) was
836 run for 120 hours. Performance benchmark inference
837 runs (A6000x1) were run over the course of
838 60 hours.

839 Each model was trained initially for 200
840 epochs¹⁶. However, the Chinese model output was
841 mostly (95%+) empty at 200 epochs and was
842 trained for another 100 epochs.

- 843 • **en-zh** byte: 296 epochs
- 844 • **en-zh** ours: 296 epochs
- 845 • **en-ja** byte: 200 epochs
- 846 • **en-ja** ours: 201 epochs
- 847 • **ja-ko** byte: 93 epochs
- 848 • **ja-ko** ours: 93 epochs

¹⁶The early stop criteria until training stalled for 10 epochs in a row.

849 A.3 Artifacts and Licensing

850 In the scope of this work, we created a reference
851 implementation and a pre-trained model as scientific
852 artifacts. The reference implementation and
853 pre-trained models will be distributed under the
854 MIT license at [To be populated after CR]. The reference
855 implementation will not contain the Llama2
856 tokenizer with our method, but an implementation
857 to create one from a downloaded Llama2 will be.

858 The Japanese-Korean task was a dataset created
859 by merging multiple datasets and could be considered
860 a novel artifact. However, as we do not have
861 redistribution rights, we will publish only the sentence
862 IDs for reproducibility purposes.

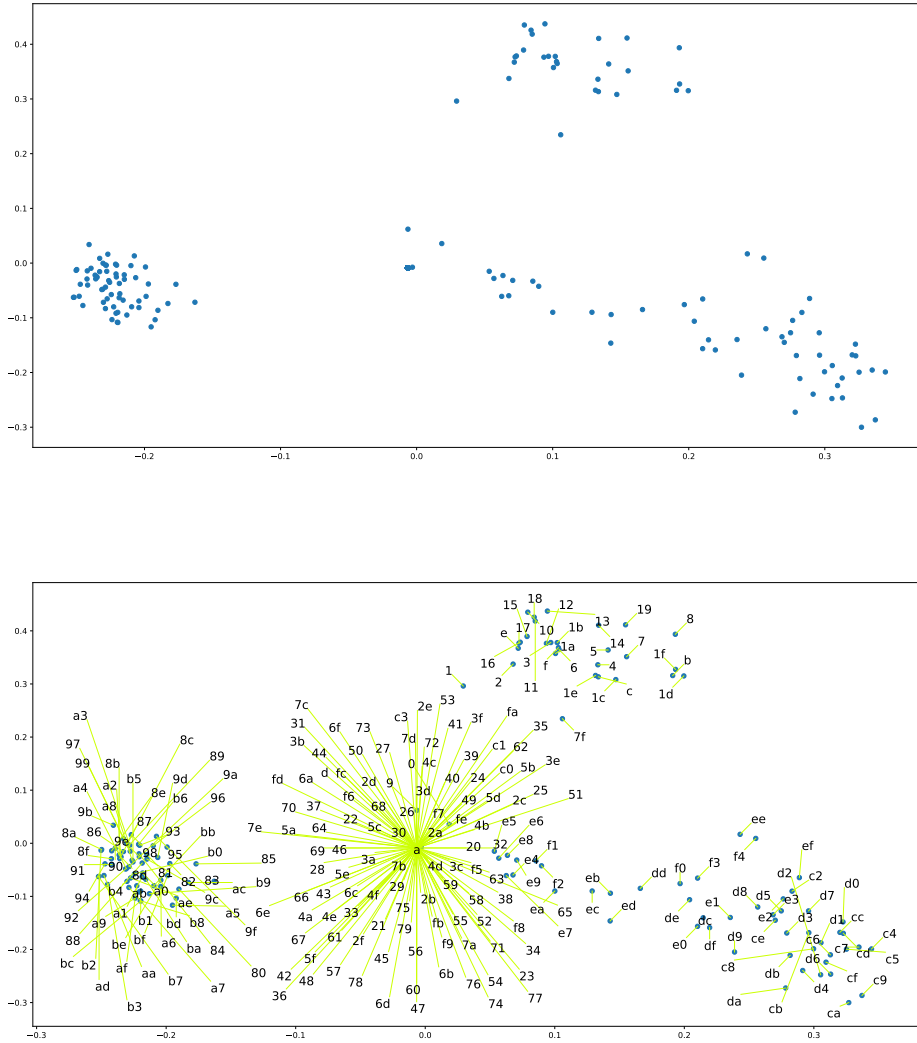


Figure 4: Byte embeddings in Llama2, dimensionality reduced with PCA ($d = 2$). We observe there are a couple clusters formed here, along with a dense cluster of undertrained tokens. The undertrained tokens in the projected space converge around the unit vector, and the number of underused tokens can be observed by the number of labels pointing to the unit vector in the lower diagram.

```
# Chinese (bleu, chrF, ter)
nrefs:1|case:mixed|eff:no|tok:zh|smooth:exp|version:2.4.2
nrefs:1|case:mixed|eff:yes|nc:6|nw:0|space:no|version:2.4.2
nrefs:1|case:lc|tok:tercom|norm:no|punct:no|asian:yes|version:2.4.2

# Japanese (bleu, chrF, ter)
nrefs:1|case:mixed|eff:no|tok:ja-mecab-0.996-IPA|smooth:exp|version:2.4.2
nrefs:1|case:mixed|eff:yes|nc:6|nw:0|space:no|version:2.4.2
nrefs:1|case:lc|tok:tercom|norm:no|punct:no|asian:yes|version:2.4.2

# Korean (bleu, chrF, ter)
nrefs:1|case:mixed|eff:no|tok:ko-mecab-0.996/ko-0.9.2-KO|smooth:exp|version:2.4.2
nrefs:1|case:mixed|eff:yes|nc:6|nw:0|space:no|version:2.4.2
nrefs:1|case:lc|tok:tercom|norm:no|punct:no|asian:yes|version:2.4.2

# Bytes (bleu, chrF)
nrefs:1|case:mixed|eff:no|tok:none|smooth:exp|version:2.4.2
nrefs:1|case:mixed|eff:yes|nc:6|nw:0|space:no|version:2.4.2
```

Table 7: sacreBLEU signatures used for score computation.