

# Egalitarian Language Representation in Language Models: It All Begins with Tokenizers

Menan Velayuthan and Kengatharaiyer Sarveswaran

Department of Computer Science, University of Jaffna, Sri Lanka.

vmenan95@gmail.com, sarves@univ.jfn.ac.lk

## Abstract

Tokenizers act as a bridge between human language and the latent space of language models, influencing how language is represented in these models. Despite the dominance of English-Centric (EC) Large Language Models (LLMs), tokenization methods often fail to fairly represent complex scripts like Tamil, Sinhala, and Hindi, primarily due to pre-tokenization choices. This study demonstrates that pre-tokenization has a more significant impact than tokenization algorithms on achieving egalitarian representation. To address this, we introduce an improvement to the Byte Pair Encoding (BPE) algorithm by incorporating graphemes, which we term Grapheme Pair Encoding (GPE). Our experiments show that grapheme-based character extraction outperforms byte-level tokenizers for complex scripts. We validate this approach through experiments on Tamil, Sinhala, and Hindi. The codebase and resources used in this work are publicly available at [GitHub Repository](#).

## 1 Introduction

Large Language Models (LLMs) have gained significant attention from the research community and the general public, especially following the public release of OpenAI’s ChatGPT in 2022 (OpenAI, 2022). LLMs have been heralded for their economic impact (Teubner et al., 2023; Eloundou et al., 2023) and have applications in areas such as coding assistants (Rozière et al., 2024; Li et al., 2023), chat systems (Meyer et al., 2023; Dorteimer et al., 2024), and machine translation (Brants et al., 2007; Dabre et al., 2020).

Mainstream LLMs are pre-trained on English-dominant corpora (Zhao et al., 2024), which makes them English-centric (EC) models. Among these, Llama 3 (Dubey et al., 2024), Mistral (Jiang et al., 2023), and Phi-3 (Abdin et al., 2024) are popular in the research community due to their open weights

and strong performance on tasks like the Massive Multitask Language Understanding (MMLU) (Hendrycks et al., 2021). This popularity has led to the use of EC LLMs as base models for developing non-English language models, such as Tamil Llama (Balachandran, 2023) and Llama 3 8B CPT SEA-LIONv2 (Lowphansirikul et al., 2021; Singapore, 2024). However, their widespread adoption often overshadows limitations in representing complex-script languages.

Petrov et al. (2024) demonstrates that the unequal treatment of languages begins at the tokenizer level. They argue that commercial LLM services charge based on token count, implying that languages requiring more tokens may disadvantage users interacting in those languages. They also discuss how tokenization differences impact latency and context window requirements, as a higher token count necessitates more computational time and larger context windows.

Our observations reveal that pre-tokenization is a key factor contributing to the unfair representation of languages with complex scripts such as Tamil, Sinhala, and Hindi. We argue that pre-tokenization limits the number of tokens a tokenizer can learn. By analyzing pre-tokenization, we can simulate the effect of a tokenizer trained with sufficient data on these languages. This helps assess whether using such a tokenizer for adapting to new complex script languages is equitable. Additionally, we show that pre-tokenization plays a more significant role than the choice of the tokenization algorithm itself (refer to Figure 1).

Byte Pair Encoding (BPE) (Sennrich et al., 2016) is a popular choice for tokenizers in EC LLMs. However, applying BPE directly to languages with complex scripts may not be optimal. To address this, we modified the BPE algorithm by using graphemes (refer to Section 3.2) as the atomic units, which we call Grapheme Pair Encoding (GPE). This adaptation enables BPE to recognize char-

acters in complex scripts similar to how humans would realize them, resulting in improved performance compared to standard BPE.

Efforts have been made to develop tokenizer-free language models (Deiseroth et al., 2024; Yu et al., 2023), which operate at the byte-level. However, this approach may not be ideal for languages with complex scripts such as Tamil, Sinhala, and Hindi. Processing these languages using graphemes proves to be more beneficial. We compare byte-level tokenizers, such as ByT5 (Xue et al., 2022) and CANINE (Clark et al., 2022), with grapheme-based character extractors and demonstrate that the latter performs better on our evaluation metrics.

## 2 Related Works

**Evaluating Tokenizers.** Goldman et al. (2024) examine the correlation between tokenizers’ text compression and language models’ performance on downstream tasks, showing a significant relationship. They conclude that compression is a reliable intrinsic indicator of tokenization quality. Rust et al. (2021) propose two metrics: 1) Normalized Sequence Length (NSL), which compares a tokenizer’s compression against a baseline like Llama, and 2) Bytes per Token, calculated by dividing UTF-8 bytes by the tokens produced. While NSL is relative, we use a variant of Bytes per Token as an absolute measure. Petrov et al. (2024) introduce *Tokenization Parity (TP)*, a metric assessing the tokenization of one language relative to another. We use *TP* to evaluate how complex scripts are underrepresented compared to English. Rust et al. (2021) also propose “subword fertility” and “proportion of continued words” as additional measures. Given their correlation with the previously discussed metrics, we choose “Compression Ratio” and “Tokenization Parity” for our evaluation.

**Tokenization Algorithms.** Senrich et al. (2016) introduce Byte Pair Encoding (BPE) as a subword segmentation strategy to manage open-vocabulary challenges in neural machine translation. This method, which improves translation accuracy by breaking down rare and unknown words into sequences of subword units, is derived from a data compression algorithm. It segments words into the most frequent pairs of bytes, facilitating a compact representation of open vocabularies using a fixed-size vocabulary of subword units. BPE-based tokenizers are a popular choice for EC Large Language

Models (LLMs). However, they may perform sub-optimally with complex scripted languages that require combinations of more than two Unicode codepoints to create characters, especially when trained with limited data and a fixed vocabulary size. Other commonly used subword tokenization techniques include WordPiece (Wu et al., 2016) and the Unigram algorithm (Kudo, 2018).

## 3 Background

This section introduces pre-tokenization and graphemes to support understanding of our research approach.

### 3.1 Pre-tokenization

Although tokenizers can be trained on large strings of textual data using Byte Pair Encoding (BPE), training them naively may lead to suboptimal performance on certain downstream tasks as these tokenizers may lead to tokens forming around common phrases or sentence (Dagan et al., 2024). To address this, it is beneficial to include a pre-tokenization step, a pre-processing step before the actual tokenization, that breaks the input text into smaller, manageable chunks. We will refer to these as *pre-tokens* throughout the paper. Since it’s inception, numerous pre-tokenization methods have been proposed, such as splitting on punctuation marks or spaces, using linguistic rule-based approaches, and regular expression-based methods (Dagan et al., 2024). HuggingFace (HF) supports various pre-tokenization techniques, and further details can be found on the pre-tokenization page<sup>1</sup>. Table 1 illustrates the pre-tokenization outputs for the translation of the text “hello!” in English, Tamil, Sinhala, and Hindi. It is evident that models like GPT-2 (Radford et al., 2019), GPT-4 (OpenAI et al., 2023), and Llama 3 (Dubey et al., 2024) unnecessarily break the text for Tamil, Sinhala, and Hindi, resulting in the need for a larger context window.

#### Pre-tokenizers based on regular expressions.

GPT-2 popularized the use of large regular expressions to segment text into smaller chunks before applying BPE. Since then, many LLMs, including closed-source models like GPT-4 and open-weight models like Llama 3, have adopted pre-tokenization to segment text before training the tokenizer. This is especially common with BPE-based tokenizers.

<sup>1</sup><https://huggingface.co/docs/tokenizers/en/api/pretokenizers>

Tokenizer	en	ta	si	hi
GPT-2	hello, !	வணக , ஃ , கம , ஃ!	ஶய , ஶு , ஶ , ஶீ , ஶி , ஶீ!	नमस , ' , त , ' !
GPT-4	hello, !	வணக , ஃகம , ஃ!	ஶய , ஶுஶ , ஶீஶி , ஶீ!	नमस , ' , त , ' !
Llama 3	hello, !	வணக , ஃகம , ஃ!	ஶய , ஶுஶ , ஶீஶி , ஶீ!	नमस , ' , त , ' !
BERT	hello, !	வணக்கம் , !	ஶயஶஶிஶி , !	नमस्ते , !
mBERT	hello, !	வணக்கம் , !	ஶயஶஶிஶி , !	नमस्ते , !
T5	hello!	வணக்கம்!	ஶயஶஶிஶி!	नमस्ते!
mT5	hello!	வணக்கம்!	ஶயஶஶிஶி!	नमस्ते!
mBART	hello!	வணக்கம்!	ஶயஶஶிஶி!	नमस्ते!
NLLBm	hello!	வணக்கம்!	ஶயஶஶிஶி!	नमस्ते!

Table 1: This table displays the outputs of the pre-tokenizer from various language models for the same word in English, Tamil, Sinhala, and Hindi.

**Pre-tokenization as tokenization bounds.** The tokenization algorithm is applied to smaller pre-tokens rather than the entire text during the tokenizer training phase. As a result, the pre-tokenization step governs the maximum possible token length that can be learned from these pre-tokens. For example, consider the text “Hello World” and assume we use pre-tokenization by splitting on whitespaces. The pre-tokens will be “Hello” and “World”. If we apply the BPE algorithm to these pre-tokens, the longest possible token learned from the first pre-token will be “Hello”. However, due to data limitations, smaller tokens may be learned from within “Hello”, such as “He”, “l”, and “lo” but the sequence cannot exceed “Hello”. Thus, the pre-tokenizer effectively bounds the maximum token length that can be formed for a given pre-token. Consequently, the Compression Ratio ( $CR$ ), as calculated using Equation 1, becomes the Maximum Compression Ratio ( $CR_{max}$ ).

### 3.2 Graphemes

Writing systems around the world vary in how they represent language, and they can be classified into six main types, namely logosyllabary/morphosyllabary, syllabary, abjad, alphabet, and abugida, based on the relationship between symbols and the spoken components of language (Daniels et al., 2003). Understanding the complexities involved in how these characters are represented in Unicode encoding is essential to processing these languages. For instance, in the abugida writing system, most of the characters are encoded using several Unicode points placed in a particular order, and when processing them, we need to treat those Unicode points together as a single unit, not separately. This sequence forms a character in

the respective language, hereafter referred to as a grapheme. For instance, “𑌕”, a Grantha grapheme encoded in Tamil Unicode, is represented using four Unicode points corresponding to the following glyphs: ‘𑌕’, ‘𑌖’, ‘𑌗’, and ‘𑌘’.<sup>2</sup> Similarly, the character ‘𑌕’ in the Sinhala language, which is also based on the abugida writing system, is a sequence of five Unicode points represented by the glyphs ‘𑌕’, ‘𑌖’, ‘𑌗’, ‘𑌘’, and ‘𑌙’, although it is considered a single character or grapheme in the language, where ‘𑌘’ is called Zero-Width Joiner. In these examples, glyphs like ‘𑌘’, ‘𑌙’, and ‘𑌚’ are called vowel modifiers, which cannot stand alone and must always be processed along with the consonant to which they are attached. These are not diacritics but vowels, represented by vowel modifiers<sup>3</sup>. These modifiers take different shapes when joined with different consonants. The pulli – ‘𑌖’ in Tamil – is not a standalone symbol but part of a pure consonant. For instance, ‘𑌕’ is the pure consonant ‘k’ in Tamil, and when a vowel (or vowel modifier), such as ‘a’, is added, the pulli will disappear. This understanding is important for Natural Language Processing (NLP) development, including tokenizers, because if an NLP tool breaks apart a vowel modifier, the result will not make much sense. Table 2 illustrates another example where how graphemes are formed in Abugida writing system based scripts.

## 4 Software and Other Specifications.

**Software Specifications.** All code for the experiments was written in Python. We utilized tokenizers from the tokenizers<sup>4</sup> library, which is part of the HuggingFace (HF) Transformers frame-

<sup>2</sup><https://www.unicode.org/charts/PDF/U0B80.pdf>

<sup>3</sup><https://www.unicode.org/charts/PDF/U0B80.pdf>

<sup>4</sup><https://github.com/huggingface/tokenizers>

	ta	si	hi
Unicode codepoint	ந, ண, ஂ, ற, ி	ස, ්, ක, ඌ, ක, ෆ, ය, ෆ	ध, न, ्र, य, व, ा, द
Grapheme Clusters	ந, ண, றி	ස්, ඞ, කි, යි	ध, न्, य, वा

Table 2: The table demonstrates how characters are perceived when separated at the Unicode codepoint and the grapheme cluster level. This comparison is performed on the phrase ‘‘Thank you’’ translated into Tamil, Sinhala, and Hindi.

work (Wolf et al., 2020). For obtaining grapheme clusters, we used the grapheme<sup>5</sup> Python library. We utilized the BPE implementation<sup>6</sup> from Karpathy (2024) to develop our code base for GPE.

**Tokenizers Used.** For our experiments, we use pre-trained tokenizers from both EC and Multi-Lingual (ML) models. EC models include GPT-2 (Radford et al., 2019), GPT-4 (OpenAI et al., 2023), Llama 3 (Dubey et al., 2024), FLAN-T5 (Chung et al., 2022), and Gemma 2 (Team et al., 2024). ML models include Aya (Üstün et al., 2024), multi-lingual BERT (Devlin et al., 2019) (referred to as mBERT), mT5 (Xue et al., 2021), mBART (Liu et al., 2020), and NLLBm (Team et al., 2022). Note that while mBERT, mT5, and mBART have EC counterparts (BERT, T5, and BART, respectively), these EC models are excluded from our analysis as they yielded identical results to their multi-lingual versions. We also utilize Byte-level tokenizers, specifically ByT5 (Xue et al., 2022) and CANINE (Clark et al., 2022). All models are available on Hugging Face (HF), except for GPT-4. We obtained the pre-tokenization regular expression for GPT-4 from (Dagan et al., 2024). We specifically use tokenizers that include a pre-tokenizer, confirmed by checking `tokenizer.backend_tokenizer.pre_tokenizer` is not None in our implementation.

**Languages for Evaluation.** We primarily focus on three South Asian languages which are based on the Abugida writing system : Hindi (hi), Tamil (ta), and Sinhala (si).

**Training and Testing Data.** For training, we randomly sample 150k Tamil sentences from the Samanantar Dataset (Ramesh et al., 2021). For testing, we use the FLORES+ (Team et al., 2022) development testsets for Tamil (ta), Hindi (hi), and Sinhala (si). We fix the vocabulary size to 5k.

<sup>5</sup><https://github.com/alvinlindstam/grapheme>

<sup>6</sup><https://github.com/karpathy/minbpe>

## 5 Methodology

In this section, we have described two separate methodologies for 1) analyzing pre-tokenization and Byte-level tokenizers, and 2) our proposed Grapheme Pair Encoding (GPE) tokenization.

### 5.1 Methodology for Analyzing Pre-tokenization and Byte-Level Tokenizers

For this analysis, we rely on the metrics **Compression Ratio** (CR) and **Tokenization Parity** (TP). We define the Compression Ratio (CR) as:

$$CR = \frac{\text{Original Sequence Length}}{\text{Tokenized Sequence Length}} \quad (1)$$

For *TP*, we adopt the definition by Petrov et al. (2024). The parity of sentence  $A$  ( $s_A$ ) relative to sentence  $B$  ( $s_B$ ) by tokenizer  $t$  is defined as:

$$TP = \frac{|t(s_A)|}{|t(s_B)|} \quad (2)$$

where  $t(s_A)$  represents the tokenization of sentence  $s_A$ , and  $|t(s_A)|$  denotes its length. It could be stated that tokenizer  $t$  achieves parity for  $A$  with respect to  $B$  when the *TP* is close to 1 (Petrov et al., 2024).

Since we consider the pre-tokenization outputs for these calculations, as explained in Section 3.1, the calculated *CR* represents the maximum *CR* ( $CR_{max}$ ), as pre-tokenization determines the maximum number of tokens present.

**Analyzing Pre-tokenization.** We evaluate English-Centric and Multi-Lingual language models based on their pre-tokenization outputs rather than the final output of the tokenizers. Tokenizer training practices, including the choice of datasets and hyperparameters, vary across development teams. Consequently, comparisons of tokenizers based solely on their final outputs can be misleading. Instead, we focus on the pre-tokenization step, as detailed in Section 3.1. The longest sequence a tokenizer can learn

in the pre-tokenization is the pre-token itself. Although smaller sub-words might be learned with insufficient training data, in an ideal case where ample data is provided, the learned tokens will match the pre-tokens. By analyzing the outputs of pre-tokenizers rather than those of the tokenizers themselves, we gain a clearer understanding of the ideal tokenization a tokenizer can achieve. This insight is crucial when assessing which tokenizers are best suited for complex script languages, particularly under the assumption that ample data will be provided when integrating a complex script language with the tokenizer. As shown in Figure 1, the impact of pre-tokenization is far more significant than that of the tokenization algorithm itself. Results of the pre-tokenization analysis are presented in Tables 3, 4, and 1.

**Analyzing Byte-level Tokenizers.** Since Byte-level tokenizers operate directly on the fundamental representation of text in digital form (bytes), analyzing them at the pre-tokenization level is not applicable. Therefore, we evaluate these tokenizers based on their final output. Results for Byte-level tokenizers are presented in Tables 5 and 6.

## 5.2 Grapheme Pair Encoding (GPE)

Instead of considering bytes as the smallest units, as done in BPE, we consider graphemes as the smallest units (refer Section 3.2). This is achieved by introducing a pre-processing step that breaks the given text into graphemes and updates the initial vocabulary with the unique graphemes present in the tokenizer training data. Once the initial vocabulary is updated, the remainder of the method adheres to the standard BPE algorithm, but operates on graphemes. Our proposed methodology is detailed in Algorithm 1. We compare the GPE approach against vanilla implementations of BPE, Unigram, and WordPiece algorithms. All tokenizers are trained on a randomly sampled 150k sentences from the Samanantar Tamil dataset (Ramesh et al., 2021) and tested on the FLORES+ Tamil development testset (Team et al., 2022). For a fair comparison, we use a whitespace-based pre-tokenizer for all algorithms.

## 6 Results and Discussion

In this section, we present and discuss the results of our experiments on pre-tokenization and the proposed GPE algorithm.

---

### Algorithm 1: Grapheme Pair Encoding (GPE)

---

**Input:** Dataset  $D$  of lines  $N$ , Vocabulary size  $|V|$ , Pre-tokenization regular expression  $RE$

**Output:** Vocabulary  $V$ , Merges  $M$

```

// init Vocabulary and Merges
 $V \leftarrow \{\}$ 
 $M \leftarrow \{\}$ 
// init storing unique graphemes
 $ghs \leftarrow \{\}$ 
for  $i \leftarrow 1$  to  $N$  do
    // extract pre-tokens based on RE
     $pre - tokens \leftarrow \text{Extract}(D[i], RE)$ 
    for each  $pt_i \in pre - tokens$  do
        // get graphemes  $gs_i$  for each  $pt_i$ 
         $gs_i \leftarrow \text{GetGraphemes}(pt_i)$  // get unique graphemes from  $gs_i$  compared to  $ghs$ 
         $gs_{unique} \leftarrow \text{GetUnique}(gs_i, ghs)$ 
        // update  $ghs$  with the new graphemes
         $ghs \leftarrow \text{Update}(gs_{unique}, ghs)$ 
    end
end
// update vocabulary with  $ghs$ 
 $V \leftarrow \text{Update}(ghs, V)$ 
// follow standard BPE merges
 $V, M \leftarrow \text{BPE}(D, V, M)$ 

```

---

## 6.1 English-Centric (EC) and Multi-Lingual (ML) Models

Tables 3 and 4 present the Maximum Compression Ratio ( $CR_{max}$ ) and  $TP$ , respectively (see Section 3.1 for details). For this sub-section, it is important to note that  $CR_{max}$  and  $TP$  are calculated based on the pre-tokenization outputs rather than the tokenizer outputs. We believe this approach provides a fair comparison across all tokenizers and establishes an upper bound for  $CR$ . The observations from both tables are discussed jointly, as they complement each other.

As expected, both EC and ML models demonstrate a strong compression ratio of  $5\times$  for the English language. This can be attributed to the simplicity of the English script; being part of the ASCII system, English characters fit into a single byte in UTF-8 representation.

Tokenizer		en	ta	si	hi
GPT-2	EC	5.26	1.36	1.55	1.56
GPT-4	EC	5.23	2.13	2.16	2.04
Llama 3	EC	5.23	2.13	2.16	2.04
FLAN-T5	EC	6.06	9.21	6.34	5.13
Gemma 2	EC	6.06	9.21	6.34	5.13
Aya	ML	6.06	9.21	6.34	5.13
mBERT	ML	5.22	7.77	5.63	4.59
mT5	ML	6.06	9.21	6.34	5.13
mBART	ML	6.06	9.21	6.34	5.13
NLLBm	ML	6.06	9.21	6.34	5.13

Table 3: The table shows the maximum compression ratio achievable for the four languages based on the pre-tokenization functions for both English Centric and Multi-Lingual models.

The GPT-2 pre-tokenizer exhibits the lowest performance, with a  $CR_{\max}$  score of only  $1.6\times$  across Tamil, Sinhala, and Hindi — all non-English languages. This indicates that the maximum compression ratio achievable by a GPT-2 tokenizer, even with sufficient training data for these languages, will not exceed  $1.6\times$ . Similarly, GPT-2 demonstrates the highest  $TP$  across all compared non-English languages, with Tamil having a  $TP$  of 4.54. This result suggests that, on average, at least 4.54 Tamil tokens are required for every English token. The  $TP$  can also be interpreted as a proxy for the relative context window size, implying that GPT-2 requires a context window 4.54 times longer for Tamil compared to English. Consequently, from a tokenization perspective, GPT-2 appears unsuitable for adapting to languages like Tamil, Sinhala, and Hindi.

Both GPT-4 and Llama 3 show better performance than GPT-2, achieving a  $CR_{\max}$  greater than  $2\times$  and a  $TP$  well below 3.0. However, despite these improvements, GPT-4 and Llama 3 still underperform when compared to the ML models, as well as the EC models FLAN-T5 and Gemma 2. We hope this observation serves as a cautionary note for researchers utilizing or adapting GPT-4 and Llama 3 for non-English languages.

As expected, ML models perform well on both the  $CR_{\max}$  and  $TP$  metrics. Interestingly, FLAN-T5 and Gemma 2 perform on par with the best-performing ML models for non-English languages. This observation suggests that FLAN-T5 and Gemma 2 could be promising models for adapting to new languages. We leave the exploration

Tokenizer		ta	si	hi
GPT-2	EC	4.54	3.41	3.38
GPT-4	EC	2.89	2.42	2.56
Llama 3	EC	2.89	2.42	2.56
FLAN-T5	EC	0.78	0.96	1.18
Gemma 2	EC	0.78	0.96	1.18
Aya	ML	0.78	0.96	1.18
mBERT	ML	0.80	0.93	1.13
mT5	ML	0.78	0.96	1.18
mBART	ML	0.78	0.96	1.18
NLLBm	ML	0.78	0.96	1.18

Table 4: The table shows the tokenization parity relative to English for the three languages in both English Centric and Multi-Lingual models. This can also be interpreted as the context window size relative to English.

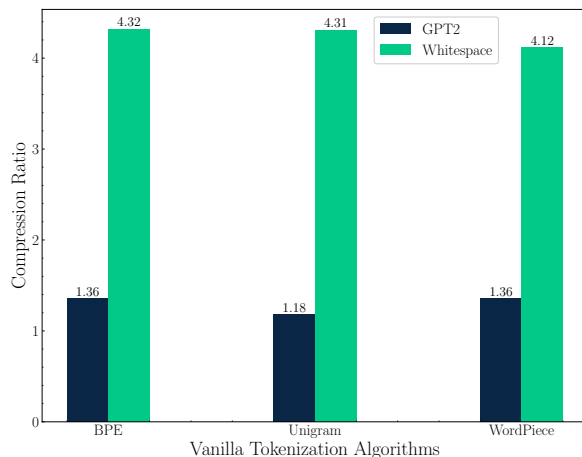


Figure 1: The figure illustrates the compression ratio of vanilla tokenizers trained on a randomly sampled 150k sentences of the Samanantar Tamil dataset using GPT-2 and whitespace-based pre-tokenization, and tested on the FLORES+ Tamil development testset.

of the impact of selecting non-English language-friendly pre-tokenization on language modeling for future work.

The analysis underscores pre-tokenization’s pivotal role in shaping language representation within models. To substantiate this claim, we trained vanilla tokenization algorithms — BPE, Unigram, and WordPiece—on a sufficiently large Tamil dataset (a randomly sampled 150k tamil sentences of the Samanantar dataset) with a vocabulary size of 5k. For this study, we utilized two pre-tokenization methods: (1) GPT-2 pre-tokenization, and (2) Whitespace pre-tokenization. Figure 1 illustrates the compression ratio of tokenizers trained

Tokenizer	en	ta	si	hi
CANINE	0.98	0.99	0.98	0.98
ByT5	0.99	0.37	0.38	0.39
Grapheme based	1.0	1.55	1.41	1.45

Table 5: Compression ratios for four languages using Byte-Level Tokenizers and our grapheme-based character extractor.

under these conditions, tested using the FLORES+ development testset.

The results of the tokenizers trained with GPT-2 pre-tokenization show nearly equal and consistently poor performance. Specifically, the BPE algorithm trained with GPT-2 pre-tokenization validates the predicted  $CR_{\max}$  score in Table 3 for GPT-2; even with a sufficiently large dataset, the compression ratio did not exceed 1.36. In contrast, tokenizers trained using simple whitespace pre-tokenization outperform those trained with GPT-2 pre-tokenization by a significant margin. Moreover, all tokenizers trained with whitespace pre-tokenization exhibit similar performance levels.

Given that the tokenizers trained with GPT-2 and whitespace pre-tokenization methods show comparable performance within their respective groups, this finding demonstrates that the compression ratio is primarily determined by the pre-tokenization methodology employed, rather than the specific tokenization algorithm used.

## 6.2 Evaluating Byte-Level Tokenization

Tables 5 and 6 present a comparison of Byte-level tokenizers and Grapheme-based character extractor based on  $CR$  and  $TP$ . We utilize two byte-level tokenizers: ByT5, which employs UTF-8 encoding to break text into bytes, and CANINE, which uses UTF-32 encoding based on Unicode codepoints. This approach allows us to handle all Unicode characters, including rare symbols and emojis, without relying on extensive vocabularies or complex pre-processing. Comparing these byte-level methods with grapheme cluster-based character-level tokenization is fair and informative, as both tokenize text at a fundamental level. Our comparison highlights the advantages of aligning tokenization with graphemes, particularly for languages with complex scripts.

In Table 5, we observe that the  $CR$  for CANINE is close to 1 for all languages, attributed to its UTF-32 representation of each character. In contrast, ByT5, which uses UTF-8 encoding,

Tokenizer	ta	si	hi
CANINE	1.17	1.0	1.0
ByT5	3.2	2.62	2.55
Grapheme based	0.76	0.71	0.69

Table 6: Tokenization Parity relative to English for three languages using Byte-Level tokenizers and our grapheme-based character extractor.

performs poorly on non-English languages such as Tamil, Sinhala, and Hindi, as these scripts require multiple bytes per character in UTF-8 encoding scheme. Grapheme-based tokenization, however, shows a significant performance improvement, with  $CR$  exceeding  $1.4\times$ . This improvement is due to graphemes effectively capturing characters by combining multiple codepoints into a single unit. Table 2 illustrates this with a comparison of how the phrase ‘‘Thank you’’ is represented at the Unicode codepoint and grapheme levels in Tamil, Sinhala, and Hindi.

Table 6 displays the  $TP$  for the Byte-level and Grapheme cluster-based tokenizers. CANINE achieves near-parity with English for Tamil, Sinhala, and Hindi, indicating it requires a similar context window size for these languages. ByT5 performs the worst among the tokenizers, with Tamil having the lowest  $TP$  of 3.2, meaning Tamil text requires three times the context window size of English. Sinhala and Hindi require at least  $2.5\times$  the context window size compared to English. In contrast, grapheme-based tokenization achieves the best performance, requiring less than 0.76 times the context window size of English for Tamil, Sinhala, and Hindi.

The findings establish grapheme-based tokenization as a superior approach for handling complex-script languages, particularly in character-based tokenization tasks. It enhances character representation by aligning more closely with human perceived characters. The impact of this choice on language modeling is a topic for future research and is beyond the scope of this paper.

## 6.3 Evaluating Grapheme Pair Encoding

We evaluate our proposed Grapheme Pair Encoding (GPE) alongside traditional tokenization algorithms such as BPE, Unigram, and WordPiece. To ensure a fair comparison, we utilize the Whitespace pre-tokenizer from HuggingFace. Details of the training process are outlined in Section 5.2.

Table 7 contains  $CR$  for Tamil from the afore-

BPE	Unigram	WordPiece	GPE
4.32	4.31	4.12	<b>4.36</b>

Table 7: Compression ratios of vanilla tokenization algorithms and Grapheme Pair Encoding, trained on a 150k Tamil sentences from the Samanantar dataset and tested on the FLORES+ Tamil development testset.

mentioned experiment. Our proposed method, GPE, better than all other tokenization algorithms, though the improvement is not significant. While GPE, being a derivative of the BPE algorithm, achieves a better  $CR$ , the difference of 0.04 is relatively minor. The WordPiece algorithm shows the poorest performance among them, but all algorithms achieve a  $CR$  greater than 4. This supports the conclusion from Figure 1 that the pre-tokenizer has a more significant impact on Compression Ratio and Tokenizer Parity (and indirectly, context window size) than the choice of tokenization algorithm.

We developed GPE by incorporating graphemes into the BPE algorithm, chosen for its straightforward implementation. In future work, we aim to integrate graphemes into other algorithms, such as Unigram and WordPiece, to further expand the scope of GPE.

## 7 Conclusion

In this work, we demonstrate the crucial role of pre-tokenization in achieving an egalitarian representation of languages in language models. To the best of our knowledge, this is the first study to focus specifically on pre-tokenizers. Our findings reveal that popular English-Centric language models inadequately represent complex scripted languages like Tamil, Sinhala, and Hindi. This underscores the need for caution when using English-Centric models as the base for developing language-specific LLMs. We advocate for the exploration of tokenization choices used by Multi-Lingual models, as they are progressing towards more egalitarian language representation.

Our analysis shows that pre-tokenization significantly affects tokenization, often more than the choice of algorithm. We propose improving BPE by incorporating graphemes, creating our Grapheme Pair Encoding (GPE) method. However, our focus is on the tokenizer, leaving broader implications for language model performance to future research. We aim for this work to inspire

further research, advancing egalitarian language representation in NLP.

## Limitations

Our work, like any other, has its limitations. We focused on the intrinsic evaluation of tokenizers rather than their impact on downstream tasks. Additionally, our study primarily explored the Abugida script, with limited attention to the Abjad writing system. Moreover, the accuracy of our results relies on the precision of the open-source grapheme library employed. These limitations represent meaningful research directions, and we hope they inspire further exploration by the research community.

## Acknowledgement

We acknowledge the support received from the LK Domain Registry in publishing this paper. This research study was carried out under the DigSAL project, supported by the German Academic Exchange Service (DAAD) and funded by the Federal Ministry for Economic Cooperation and Development (BMZ) through SDG Partnerships.

We extend our gratitude to Dr. Andrej Karpathy for his educational lecture<sup>7</sup> on Byte Pair Encoding, available on YouTube, which served as an inspiration for this research.

## References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, ..., and Xiren Zhou. 2024. [Phi-3 Technical Report: A Highly Capable Language Model Locally on Your Phone](#). *Preprint*, arXiv:2404.14219.
- Abhinand Balachandran. 2023. [Tamil-Llama: A New Tamil Language Model Based on Llama 2](#). *Preprint*, arXiv:2311.05845.
- Thorsten Brants, Ashok C. Popat, Peng Xu, Franz J. Och, and Jeffrey Dean. 2007. [Large Language Models in Machine Translation](#). In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pages 858–867, Prague, Czech Republic. Association for Computational Linguistics.

<sup>7</sup><https://youtu.be/zduSFxRajkE?si=PSWMcRiDoXxSsd2Z>



- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, ..., and Jason Wei. 2022. [Scaling Instruction-Finetuned Language Models](#). *Preprint*, arXiv:2210.11416.
- Jonathan H. Clark, Dan Garrette, Iulia Turc, and John Wieting. 2022. [Canine: Pre-training an Efficient Tokenization-Free Encoder for Language Representation](#). *Transactions of the Association for Computational Linguistics*, 10:73–91.
- Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. [A Survey of Multilingual Neural Machine Translation](#). *ACM Comput. Surv.*, 53(5).
- Gautier Dagan, Gabriel Synnaeve, and Baptiste Rozière. 2024. [Getting the most out of your tokenizer for pre-training and domain adaptation](#). *Preprint*, arXiv:2402.01035.
- Peter T Daniels, Mark Aronoff, and Janie Rees-Miller. 2003. *Writing Systems*, page 43–80. Blackwell.
- Björn Deiseroth, Manuel Brack, Patrick Schramowski, Kristian Kersting, and Samuel Weinbach. 2024. [T-FREE: Tokenizer-Free Generative LLMs via Sparse Representations for Memory-Efficient Embeddings](#). *Preprint*, arXiv:2406.19223.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). *Preprint*, arXiv:1810.04805.
- Jonathan Dornheimer, Nik Martelaro, Aaron Sprecher, and Gerhard Schubert. 2024. [Evaluating large-language-model chatbots to engage communities in large-scale design projects](#). *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 38:e4.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, ..., and Zhiwei Zhao. 2024. [The Llama 3 Herd of Models](#). *Preprint*, arXiv:2407.21783.
- Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. [GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models](#). *Preprint*, arXiv:2303.10130.
- 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](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 2274–2286, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Massive Multitask Language Understanding](#). *Preprint*, arXiv:2009.03300.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7B](#). *Preprint*, arXiv:2310.06825.
- Andrej Karpathy. 2024. minbpe. <https://github.com/karpathy/minbpe>.
- Taku Kudo. 2018. [Subword Regularization: Improving Neural Network Translation Models with Multiple Subword Candidates](#). *Preprint*, arXiv:1804.10959.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, ..., and Harm de Vries. 2023. [StarCoder: may the source be with you!](#) *Preprint*, arXiv:2305.06161.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual Denoising Pre-training for Neural Machine Translation](#). *Preprint*, arXiv:2001.08210.
- Lalita Lowphansirikul, Charin Polpanumas, Nawat Jantrakulchai, and Sarana Nutanong. 2021. [WangchanBERTa: Pretraining transformer-based Thai Language Models](#). *Preprint*, arXiv:2101.09635.
- Jesse G. Meyer, Ryan J. Urbanowicz, Patrick C. N. Martin, Karen O’Connor, Ruowang Li, Pei-Chen Peng, Tiffani J. Bright, Nicholas Tatonetti, Kyoung Jae Won, Graciela Gonzalez-Hernandez, and Jason H. Moore. 2023. [ChatGPT and large language models in academia: opportunities and challenges](#). *BioData Mining*, 16(1):20.
- OpenAI. 2022. [ChatGPT: Optimizing Language Models for Dialogue](#). Accessed: 2024-08-15.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, ..., and Barret Zoph. 2023. [GPT-4 Technical Report](#). *Preprint*, arXiv:2303.08774.
- Aleksandar Petrov, Emanuele La Malfa, Philip Torr, and Adel Bibi. 2024. Language model tokenizers introduce unfairness between languages. *Advances in Neural Information Processing Systems*, 36.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. [Language models are unsupervised multitask learners](#). *OpenAI blog*, 1(8):9.
- Gowtham Ramesh, Sumanth Doddapaneni, Aravindh Bheemaraj, Mayank Jobanputra, Raghavan AK, Ajitesh Sharma, Sujit Sahoo, Harshita Diddee, Mahalakshmi J, Divyanshu Kakwani, Navneet Kumar, Aswin Pradeep, Srihari Nagaraj, Kumar Deepak, Vivek Raghavan, Anoop Kunchukuttan, Pratyush Kumar, and Mitesh Shantadevi Khapra. 2021. [Samanantar: The Largest Publicly Available Parallel Corpora Collection for 11 Indic Languages](#). *Preprint*, arXiv:2104.05596.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Romain Sauvestre, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, ..., and Gabriel Synnaeve. 2024. [Code Llama: Open Foundation Models for Code](#). *Preprint*, arXiv:2308.12950.
- 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 *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 3118–3135, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural Machine Translation of Rare Words with Subword Units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- AI Singapore. 2024. SEA-LION (Southeast Asian Languages In One Network): A Family of Large Language Models for Southeast Asia. <https://github.com/aisingapore/sealion>.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, ..., and Alek Andreev. 2024. [Gemma 2: Improving Open Language Models at a Practical Size](#). *Preprint*, arXiv:2408.00118.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, ..., and Jeff Wang. 2022. [No Language Left Behind: Scaling Human-Centered Machine Translation](#). *Preprint*, arXiv:2207.04672.
- Timm Teubner, Christoph M. Flath, Christof Weinhardt, Wil van der Aalst, and Oliver Hinz. 2023. [Welcome to the Era of ChatGPT et al](#). *Business & Information Systems Engineering*, 65(2):95–101.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. Rush. 2020. [Transformers: State-of-the-Art Natural Language Processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Łukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, ..., and Jeffrey Dean. 2016. [Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation](#). *Preprint*, arXiv:1609.08144.
- 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](#). *Preprint*, arXiv:2105.13626.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). *Preprint*, arXiv:2010.11934.
- Lili Yu, Dániel Simig, Colin Flaherty, Armen Aghajanyan, Luke Zettlemoyer, and Mike Lewis. 2023. [MEGABYTE: Predicting Million-byte Sequences with Multiscale Transformers](#). *Preprint*, arXiv:2305.07185.
- Jun Zhao, Zhihao Zhang, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024. [Llama beyond english: An empirical study on language capability transfer](#). *Preprint*, arXiv:2401.01055.
- Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya Model: An Instruction Finetuned Open-Access Multilingual Language Model](#). *Preprint*, arXiv:2402.07827.