

L3TC: Leveraging RWKV for Learned Lossless Low-Complexity Text Compression

Junxuan Zhang^{2*}, Zhengxue Cheng^{1*†}, Yan Zhao¹, Shihao Wang²,
Dajiang Zhou², Guo Lu¹, Li Song¹

¹Institute of Image Communication and Network Engineering, Shanghai Jiao Tong University

²Ant Group, Hangzhou, China

junxuan.zjx@antgroup.com, zxcheng@sjtu.edu.cn, zhaoyanzy@sjtu.edu.cn, shihao.wsh@antgroup.com,
dajiang.zdj@antgroup.com, luguo2014@sjtu.edu.cn, song_li@sjtu.edu.cn

Abstract

Learning-based probabilistic models can be combined with an entropy coder for data compression. However, due to the high complexity of learning-based models, their practical application as text compressors has been largely overlooked. To address this issue, our work focuses on a low-complexity design while maintaining compression performance. We introduce a novel Learned Lossless Low-complexity Text Compression method (L3TC). Specifically, we conduct extensive experiments demonstrating that RWKV models achieve the fastest decoding speed with a moderate compression ratio, making it the most suitable backbone for our method. Second, we propose an outlier-aware tokenizer that uses a limited vocabulary to cover frequent tokens while allowing outliers to bypass the prediction and encoding. Third, we propose a novel high-rank reparameterization strategy that enhances the learning capability during training without increasing complexity during inference. Experimental results validate that our method achieves 48% bit saving compared to *gzip* compressor. Besides, *L3TC* offers compression performance comparable to other learned compressors, with a $50\times$ reduction in model parameters. More importantly, *L3TC* is the fastest among all learned compressors, providing real-time decoding speeds up to megabytes per second.

Code —

<https://github.com/alipay/L3TC-leveraging-rwkv-for-learned-lossless-low-complexity-text-compression.git>

1 Introduction

Lossless text compression is a fundamental research field focused on reducing data size based on information theory. In 1948, Shannon (Shannon. 1948) described that maximizing the \log_2 -likelihood of data is equivalent to minimizing the number of bits required for compression. Then information theory (MacKay. 2003) established the essential equivalence between probabilistic models of data and lossless compression. Recently, learning-based probabilistic models, like language models, have developed rapidly

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Compression Ratio vs. Model Size: The proposed *L3TC* method outperforms other well-known compressors, including *gzip* and learning-based models, with the best compression ratio and a $50\times$ reduction in model size while maintaining comparable performance. When running on devices, other learned models typically decode at KB/s speeds, while *L3TC* achieves decoding speeds up to MB/s.

and demonstrated remarkable in-context learning capabilities across various tasks by predicting the probability of the next token. This predictive capability can be inherently combined with entropy encoding to serve as a data compressor. Entropy encoding can be implemented in various ways, including Huffman Coding (Huffman 1952), Arithmetic Coding (AC) (Pasco. 1977; Howard and Vitter 1991), and Asymmetric Numeral Systems (ANS) (Duda. 2013).

Despite the success of language models in many intelligent tasks, the literature has largely overlooked their application as a practical lossless text compressor. This is not surprising, because the decoding complexity is nearly equivalent to the model’s inference complexity, which is significantly slower compared to popular engineered compressors such as *gzip* (Pasco. 1996). For a practical usage, the data compression task basically demands a very critical real-time decoding performance, up to megabytes per second. Therefore, it is a very promising and challenging research direction to develop a low-complexity text compression algo-

rithm based on the learning-based likelihood prediction.

Learned probabilistic models using transformers and LSTMs have proven effective for lossless text compression. Specifically, nncp v3.2 (Bellard 2021), which is based on transformers, and cmix v20 (Knoll 2023), which utilizes LSTMs, achieve the top-2 compression efficiencies on the enwiki9 dataset (Mahoney 2024). However, they both involve significant decoding complexity, taking 2.8 days and 7.2 days to decode 1GB of text, respectively. Recently, LLMZip (Valmeekam et al. 2023) demonstrates that using LLaMa-7B (Touvron et al. 2023a) with Arithmetic Coding can achieve about 40% bit rate savings compared to PAQ8H (Mahoney 2006a) on text8 dataset (Mahoney 2006b), but the complexity is extremely high. Additionally, Google (Delétang et al. 2024) shows that Chinchilla-70B (Delétang et al. 2024) achieves more promising bit rate savings compared to gzip on enwiki9 dataset (HutterPrize 2006). However, these models’ sizes far exceed the data to be compressed, making them impractical for real-world use.

To address the aforementioned high-complexity limitations, we explore a Learned Lossless Low-Complexity Text Compression (*L3TC*) method. It leverages the probability prediction of RWKV models while enhancing compression performance with two proposed novel designs. 1) We extensively evaluate the compression capability of various architectures, including Transformer, Transformer-XL, and RWKV, finding that RWKV offers the fastest decoding speed with moderate compression ratios. 2) We propose an outlier-aware tokenizer, focuses on frequent tokens while allowing unknown tokens to bypass encoding. 3) We introduce a novel high-rank reparameterization strategy to improve the learning capability during training without increasing inference complexity. Experiments show that *L3TC* achieves a significant 48% bit saving compared to gzip, and offers the fastest decoding speed among all the learned compressors.

The contributions of our work are summarized as follows:

- We extensively evaluate various architectures and select RWKV as our backbone, owing to its acceptable compression ratios with the fastest speed.
- We propose an outlier-aware tokenizer and a high-rank reparameterization strategy to improve the compression performance while maintaining a low-complexity design.
- Experimental results validate that *L3TC* achieves a 48% bit saving compared to gzip, reduces the model size by over 50× compared to other learned compressors with similar compression ratios, and offers decoding speeds up to the range of MB/s on mobile devices.

2 Related Work

Classical Text Compressors. The development of lossless text compression has a long history. Typical compression tools include gzip, bzip2, and zstd. gzip (Pasco. 1996) is based on the Deflate algorithm, which is a combination of LZ77 (Ziv 1977) and Huffman coding (Huffman 1952). It works by searching for duplicate strings within a sliding window and replacing them with a pointer to their previous occurrence and the string length. BZIP2 (Seward 2000) is

another popular lossless text compressor that uses Burrows-Wheeler Transform and run-length encoding, Huffman Coding. Owing to its transform to rearrange the input text into runs of similar characters, bzip2 achieves a higher compression ratio than gzip. zstd (Meta. 2015), a more recent algorithm, introduces several advanced techniques such as optimal parsing and dictionary compression, significantly improving the compression performance.

Learned Text Compressors. Transformer (Vaswani et al. 2017) or LSTM (Hochreiter and Schmidhuber 1997) based compression with arithmetic coding has achieved SOTA compression performance (Mahoney 2024). Notable examples include cmix v20 and nncp. Cmix combines multiple modeling techniques, like context mixing, prediction by partial matching, and neural network-based models, to enhance compression. Nncp uses an online compression approach, where a pseudo-randomly initialized model is continuously trained on the data stream, adapting its weights to enhance predictions over time. Language models, with their strong in-context learning abilities, are well-suited for offline compression. For instance, LLMZip, an offline compression method using pretrained LLaMa-7B, achieves up to 0.71 *bpc* (bit per character) on the text8 dataset, outperforming popular archive tools like PAQ8H (1.2**bpc**). However, these learned compressors are highly resource-intensive, requiring several days to compress 1GB of text.

Language Models and Compression. Recent studies explore the connection between compression and likelihood maximization in language models. (Huang et al. 2024) examines the linear relationship between compression and intelligence across tasks. Google (Delétang et al. 2024) demonstrates that Chinchilla-70B, pretrained on text datasets, not only achieves a promising compression ratio on text compression, but also compresses ImageNet (Deng et al. 2009) patches to 48% and LibriSpeech (Panayotov et al. 2015) audios to 21% of their raw size, obviously outperforming the domain-specific compressors like PNG (60%) (Boutell 1997) and FLAC (30%) (Xiph.Org Foundation 2001). Some research investigates the relationship between tokenization and compression. (Lester et al. 2024) introduces a novel text segmentation technique outperforming byte-level segmentation in perplexity and inference speed. Pathpiece (Schmidt et al. 2024) argues that fewer tokens do not always result in better downstream performance. (Goldman et al. 2024) suggests that compression is a reliable intrinsic indicator of tokenization quality.

Though learning-based compression offers promising results, the high complexity of billion-parameter models remains a challenge. Thus, we prioritize a low-complexity design while maintaining competitive compression ratios.

3 Proposed Method

3.1 Overall Architecture

The overall architecture of our proposed method *L3TC* is shown in Fig. 2. The text is firstly segmented into a sequence of tokens $x_{1:N}$, where N denotes the total number of tokens. Then x_i is fed into the models to calculate the probability

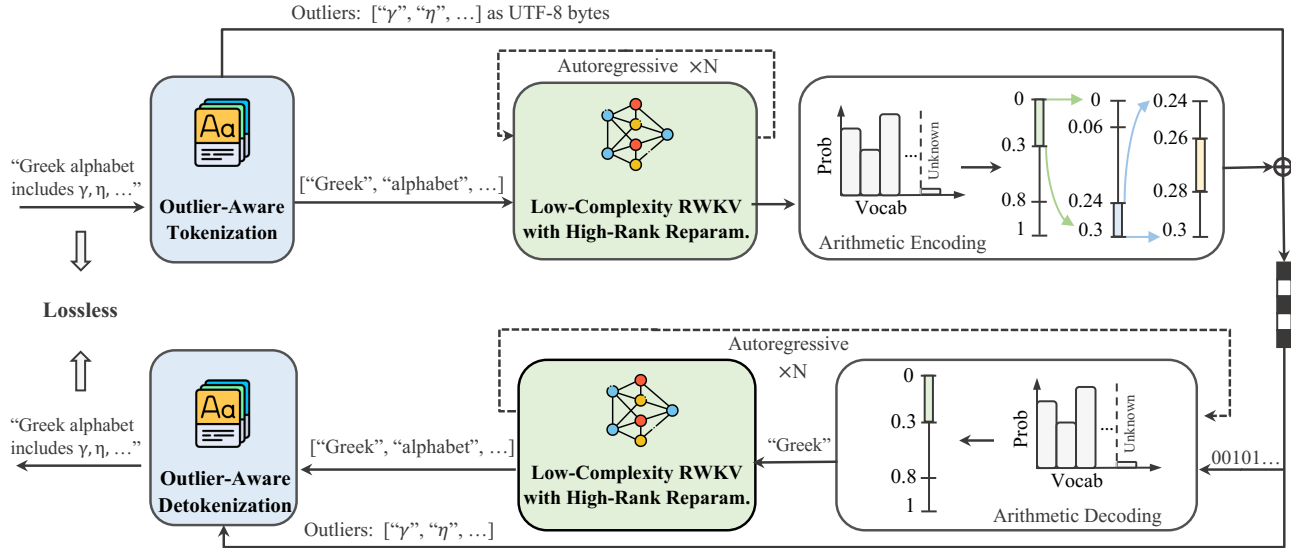


Figure 2: The proposed Learned Lossless Low-Complexity Text Compression (L3TC) framework comprises three key components: tokenization, prediction, and encoding. The text is firstly segmented through a novel outlier-aware tokenizer. Tokens in the vocabulary are then predicted by a low-complexity RWKV model and subsequently encoded by an arithmetic coder. Outliers, which appear infrequently, are allowed to bypass the prediction and encoding. A high-rank reparameterization strategy is introduced to enhance the RWKV models’s prediction capability during training without increasing inference complexity.

based on the context information, i.e., $p(x_i|x_{\leq i})$. Arithmetic coding, known for achieving optimal coding efficiency, is used to compress the data to its entropy limit. The expected (optimal) bit number of the compressed data is the entropy:

$$H(p) = \mathbb{E}\left(\sum_{i=1}^n -\log_2 p(x_i|x_{\leq i})\right) \quad (1)$$

Language models are trained using cross-entropy loss to maximize the likelihood, which is consistent with compression objectives. At the receiver side, tokens are decoded autoregressively using the probabilistic model and an arithmetic decoder as Fig. 2. As arithmetic coding complexity is negligible compared to model inference, the overall decoding complexity $\mathcal{O}(x_{1:N})$ is proportional to the number of tokens N and the network inference time $\mathcal{O}(x_i)$:

$$\mathcal{O}(x_{1:N}) = \mathcal{O}(x_i) \times N \quad (2)$$

To reduce the overall decoding complexity, we first explore a low-complexity RWKV model to minimize per-token inference time $\mathcal{O}(x_i)$ in Section 3.2. Then, we introduce an outlier-aware tokenizer that tokenizes only frequent characters and bypasses encoding for outliers, as detailed in Section 3.3. Finally, we present a high-rank reparameterization strategy in Section 3.4 to enhance training without increasing inference complexity. All these innovations form our learned lossless low-complexity text compressor (L3TC).

3.2 Low-Complexity RWKV Models

To explore a low-complexity design, we conduct experiments on various architectures, including Transformer (Vaswani et al. 2017), Transformer-XL (Dai et al.

Compressor	MACs↓	CR(%)↓	ACR(%)↓	Time↓
Tsf-200K	107.78K	31.86	31.86	8.61ms
Tsf-800K	412.16K	27.58	27.74	16.71ms
Tsf-169M	85.11M	14.81	31.71	-
TsfXL-200K	155.50K	27.14	27.14	4.36ms
TsfXL-800K	483.25K	21.49	21.65	7.57ms
TsfXL-169M	92.26M	14.81	31.71	-
RWKV-200K	143.62K	24.36	24.36	0.34ms
RWKV-800K	471.68K	21.73	21.89	0.41ms
RWKV-169M	92.19M	15.88	32.78	-

Table 1: Compression performance using Transformer (Tsf), Transformer-XL(TsfXL) and RWKV on enwiki9 dataset.

2019), and RWKV (Peng et al. 2023). Following the settings in (Delétang et al. 2024), we train models from scratch on enwik8 (HutterPrize 2006) and test on enwik9, using a character-based tokenizer and vocabulary size of 128. The prediction capability of language model is largely influenced by context length and too long context increase the running time, so we chunk the data into sequences of C bytes (C is 2048) similar to (Delétang et al. 2024).

To evaluate compression performance, we report both compression ratio (CR) and adjusted compression ratio (ACR). CR is the ratio of the compressed data size to the raw data size, while ACR accounts for model size by adding it to the compressed size, with model size calculated in float16 precision. To assess complexity, we list the Multiply-Accumulate Operations (MACs) and inference time per token on an iPhone12 ANE with a batch size of 128. For mod-

els with 169M parameters, real-time inference is unachievable, so their running times are omitted.

Some key findings can be observed from Table. 1: 1) *Compression Ratio*: Larger models generally achieve better compression, but RWKV scales slightly worse than Transformer and Transformer-XL, likely because of its limitation of linear attention. Transformer-XL and RWKV are more effective when the model size is relatively small, probably because they both propose explicit memory mechanisms to enlarge the context length. 2) *Complexity*: Transformer-XL and RWKV have higher MACs than vanilla Transformers due to their explicit memory designs. However, they benefit from the fixed context lengths, enabling faster inference. Transformers typically process variable sequence lengths, leading to slower inference with longer inputs. RWKV achieves even faster speeds with its memory length of 1 compared to Transformer-XL’s memory length of 256.

In summary, RWKV offers the fastest inference speeds with acceptable compression ratios, making it most suitable for our low-complexity design.

3.3 Outlier-aware Tokenizer

Tokenizers segment text into subword tokens, serving as a pre-compression process (Lester et al. 2024). By concealing the character-level composition of each token and merging the most frequent subwords, tokenizers enable the network to model long-distance dependencies and process more data. This approach also enhances the context prediction during inference. Consequently, enhancing tokenizers’ performance not only improves the overall compression ratio but also reduces inference complexity.

Vocabulary size and coverage are key for tokenization. Recent language models typically use vocabulary sizes ranging from 32K to 256K tokens. However, a larger vocabulary would increase the inference burden, especially for very small models. While language models usually aim for 100% character coverage, rare tokens can actually increase the compressed size and reduce the compression efficiency.

To investigate the influence of vocabulary sizes and coverage values, we conduct experiments on enwik8 using Byte Pair Encoding (BPE) (Sennrich, Haddow, and Birch 2016) as the baseline, which is recognized for its optimal pre-compression performance. We use two distribution methods, “Unigram” and “Uniform”, to determine the number of bits for each token after tokenization. “Unigram” calculates the entropy of tokens based on their contextual probabilities in enwik8, assigning shorter codes to more frequent tokens, making it suitable for evaluating in-distribution performance. “Uniform” assigns identical bits to all tokens, hence is suitable to assess out-of-distribution performance. If the character coverage is lower than 1, there exists a vocabulary set Γ containing all known tokens. Text characters excluded by the vocabulary, i.e., the unknown tokens, are identified as outliers and transmitted through a bypass. Therefore, the total compressed size *Bits* is defined as

$$Bits = H_{x_i \in \Gamma}(p) + R(x_i \notin \Gamma) \quad (3)$$

where $H(p)$ represents the compressed size of known tokens, as defined in Eq.(1), and $R(x_i \notin \Gamma)$ denote the con-

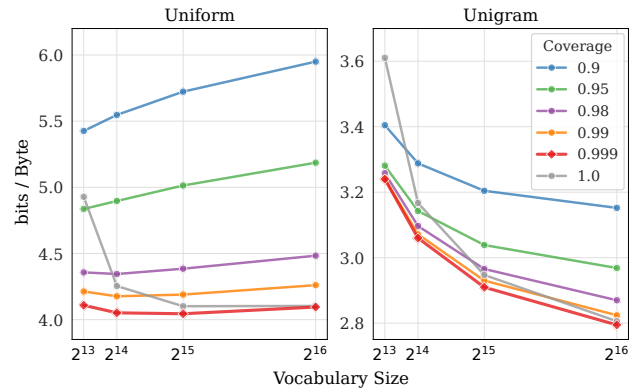


Figure 3: Performance with different coverage values.

Vocabulary Size	8K	16K	32K	64K
NPU (bs=16)	0.36ms	0.51ms	2.44ms	3.96ms
CPU (bs=1)	0.12ms	0.23ms	0.52ms	1.01ms

Table 2: Decoding speeds using different vocabulary sizes.

sumed bits for outliers to be coded directly as a UTF-8 byte sequence. Two parts contribute to their total bit sizes. The “bits per byte (bpb)” metric (the number of bits to encode one byte of text) is then calculated to assess the tokenizer’s pre-compression efficiency, with a lower value indicating better performance, as shown in Fig. 3.

It can be observed the coverage of 0.999 achieves the lowest bpb for both “Uniform” and “Unigram” methods. When using the “Unigram” method, increasing the vocabulary size can reduce the sequence length, packing more information into the context and resulting in a better compression ratio. However, this benefit diminishes if the evaluation data’s distribution differs from the training data, as measured using the “Uniform” method. In other words, the largest vocabulary size is not always optimal. Moreover, our evaluation using RWKV-800K (Table 2) shows that the inference time increases with the vocabulary size. Thus, a relatively small vocabulary size is essential for a low-complexity design.

3.4 High-rank Reparameterization

To enhance the inference performance of RWKV without compromising the computational efficiency, we introduce a High-rank (HiRA) reparameterization strategy. As shown in Fig. 4, we add additional branches for each R, K, V layer in both “Channel Mixing” and “Time Mixing” modules during training to enhance the training effectiveness, and then merge them into the main branches during inference to reduce the number of parameters. Inspired by Low-rank adaption (LoRA) (Hu et al. 2021), we increase the parameters of these additional branches through matrix decomposition, thereby improving their capability.

Specifically, assuming the original parameters of the R, K, V layers are $W_0 \in \mathbb{R}^{d \times k}$, we do not directly add several parallel branches, such as 1x1 convolution or shortcuts as in (Ding et al. 2021). Instead, given the m -th branch, we

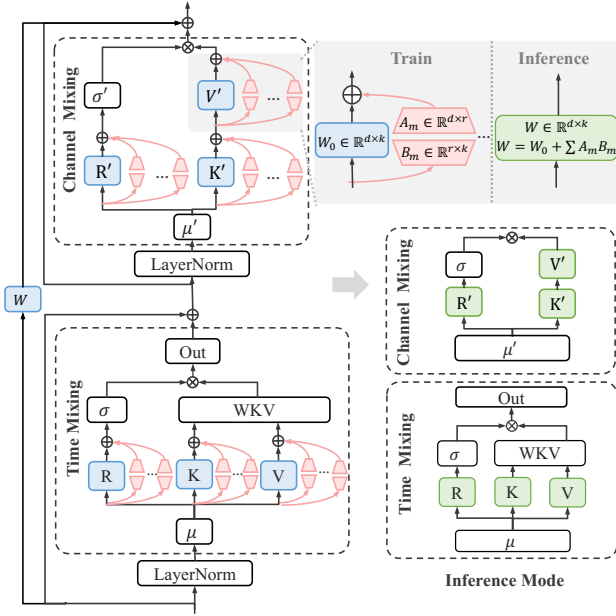


Figure 4: Proposed high-rank reparameterization method.

decompose it into the product of two high-rank matrices: $A_m \in \mathbb{R}^{d \times r}$ and $B_m \in \mathbb{R}^{r \times k}$. Therefore, during training, these layers' output is obtained by summing the outputs of the main branch and several bypass branches. Unlike LoRA, the matrices W_0 , A_m , and B_m are optimized simultaneously. The effect of the rank r and the number of branches m will be discussed in the ablation studies.

As illustrated in the right part of Fig. 4, during inference, the additional branches $A_m \times B_m$ are merged into the main branch via structural reparameterization:

$$W = W_0 + \sum_m A_m \times B_m \quad (4)$$

The reparameterized single-path structure reduces inference time and memory usage while retaining multi-branch parameters for high performance. Besides, a linear layer (depict as W) is added to the RWKV module, boosting performance with minimal computational overhead.

4 Experiments

4.1 Experimental Setup

Implementation Details. For RWKV models, we adjust the number of layers, attention embedding dimension and hidden sizes to achieve target model sizes. We train the models using the AdamW (Loshchilov and Hutter 2019) optimizer with an initial learning rate of $1e-4$ and a linear learning rate scheduler with a decay rate of 0.999 over 20 epochs, without warm-up. All the models are trained with a sequence length of 2048 bytes and a batch size of 64.

For evaluation, learned compressors process data in 2048-byte sequences with parallel batch processing, while traditional compressors with longer context lengths use unchunked data to preserve optimal compression ratios.

Dataset and Metrics. The *enwik8/enwik9* datasets, consisting of the first 1GB and 100MB of the English Wikipedia XML, are commonly used to evaluate text compression performance. Since *enwik8* only contains 10% data of *enwik9*, they represent a significant distribution shift. Therefore, we train our L3TC models on *enwik8* and evaluate them on both *enwik8* and *enwik9* to assess the in-distribution and out-of-distribution compression performance.

We evaluate compression performance using CR and ACR metrics, similar to Section 3.2, and assess complexity through MACs and decoding speeds. The decoding speeds are measured on typical computing platforms, including server GPUs (NVIDIA A100 80 GB) and device NPUs (iPhone12 Apple Neural Engine). To measure on-device performance, we convert models to CoreML packages (Apple 2023) and use Xcode software to get the inference time.

Baselines. Comparisons are conducted with classical compressors including gzip (Pasco. 1996), bzip2 (Seward 2000), zstd (Meta. 2015), and learned compressors based on pretrained language models, such as Llama2 (Touvron et al. 2023b), Chinchilla (Hoffmann et al. 2022), and RWKV (Peng et al. 2023). Models with various sizes are downloaded from Huggingface. We omit LLMZip (Valmeekam et al. 2023) as its results are equivalent to Llama-7B. Models in *tszip* (Bellard 2023) are also included in comparisons. Since Chinchilla models are not open-source, their experimental results are derived from the paper (Delétang et al. 2024). All other models' results are recorded based on our implementations. We also compare the results of transformer in (Delétang et al. 2024), which are slightly different from Table. 1, because the network structures and tokenizers are not exactly the same.

4.2 Compression Performance

Table 3 illustrates the compression ratios for different compressors on *enwik9*. Models labeled with * are pretrained models on larger text datasets using their own tokenizers. Models without * are trained from scratch on *enwiki8*. Our proposed L3TC models are highlighted in bold and in italics.

As shown in Table 3, classical compressors, e.g., gzip, offer moderate compression ratios with minimal computational complexity. Compressors based on pretrained models (Llama2, Chinchilla, and RWKV) achieve superior compression performance, with Llama2-13B leading by achieving a 75% bit saving compared to gzip. However, these models suffer from significant model sizes (ranging from 169M to 70B parameters) and excessive resource requirements. In fact, using such large models to compress small-scale data is highly inefficient, with ACR values ranging from 48% to 14,000%. Even compressed using low-bit quantization, the model size overhead remains substantial, making them impractical for real-world deployments.

By contrast, the proposed *L3TC* models have smaller model sizes and lower computational complexity. *L3TC-3.2M* achieves 50% bit saving compared to gzip and outperforms other learned compressors by saving 48% more bits when accounting for model size. *L3TC* also offers compression ratios comparable to other learned compressors, with a

Chunk	Compressor	Tokenizer	MACs	CR(%)	ACR(%)
default	gzip	-	-	32.26	32.26
	bzip2	-	-	25.40	25.40
	zstd (-22 -ultra)	-	-	21.49	21.62
	tszip* (Bellard 2023)	RWKV-Pile-Tokenizer-50K	130.71M	13.54	47.34
2048	Llama2-7B*	Llama-tiktoken-32K	6.61G	8.62	1408.62
	Llama2-13B*	Llama-tiktoken-32K	12.85G	8.04 (-75%)	2608.04
	Chinchilla-1B* (Delétang et al. 2024)	Chinchilla-tokenizer	-	11.30	211.30
	Chinchilla-7B* (Delétang et al. 2024)	Chinchilla-tokenizer	-	10.20	1410.20
	Chinchilla-70B* (Delétang et al. 2024)	Chinchilla-tokenizer	-	8.30	14008.30
	RWKV-1.5B*	RWKV-Pile-tokenizer-50K	1.41G	10.89	310.89
	RWKV-7B*	RWKV-Pile-tokenizer-50K	7.19G	9.68	1409.68
	Tsf-200K (Delétang et al. 2024)	ASCII	-	30.90	30.90
	Tsf-800K (Delétang et al. 2024)	ASCII	-	21.70	21.86
	Tsf-3.2M (Delétang et al. 2024)	ASCII	-	17.00	17.64
	L3TC-200K	Outlier-aware tokenizer	1.72M	18.48 (-43%)	18.48 (-43%)
	L3TC-800K	Outlier-aware tokenizer	3.33M	17.56 (-46%)	17.72 (-45%)
	L3TC-3.2M	Outlier-aware tokenizer	6.17M	16.23 (-50%)	16.87 (-48%)
	L3TC-12M	Outlier-aware tokenizer	12.99M	16.00 (-50%)	18.40 (-43%)

Table 3: Compression ratios on enwiki9 dataset (1GB): Bold indicates best performance. Although pretrained Llama2-13B achieves the best performance with the lowest CR, it is impractical to use. When considering the overhead of model size, our proposed *L3TC-3.2M* achieves the best performance with the lowest ACR, saving 48% bits when using *gzip* as an anchor.

Chunk	Compressor	CR(%)	ACR(%)
default	gzip	36.45	36.45
	tszip (Bellard 2023)	13.83	351.83
2048	L3TC-200K	17.55 (-52%)	17.95 (-51%)
	L3TC-800K	15.68 (-57%)	17.28 (-53%)
	L3TC-3.2M	13.81 (-62%)	20.23 (-44%)
	L3TC-12M	11.24 (-69%)	35.24 (-3%)

Table 4: Compression ratios on enwik8 dataset (100MB).

50× reduction in model parameters, as shown in Fig. 1.

Table 4 records the compression ratios on enwik8, providing an assessment of in-distribution performance. While larger model sizes generally enhance compression, L3TC-12M shows no further gains on enwik9 compared to L3TC-3.2M. Actually, when the model size approaches the size of training data, the model is likely to overfit to the training set. Hence L3TC-12M is excluded in the following sections.

4.3 Decoding Speed

Table 5 compares the decoding speeds for various methods. To calculate the decoding speed (bytes per second), we multiply the batch size by the average byte length per token and then divide the result by the batch inference time. The average byte lengths for the character-based tokenizer, Llama tiktokenizer, RWKV-pile tokenizer, and our proposed outlier-aware tokenizer are 1, 2.94, 3.37 and 3.29, respectively. On the A100, the batch size is 2048 by default or the maximum size without out-of-memory. When running on mobile devices, the batch size is set to 256. Pretrained models with hundreds of millions to billions of parameters are difficult to run on devices with limited memory and compu-

Compressor	iPhone12@ANE	GPU@A100
Llama2-7B* (LLMZip)	-	280 B/s
Llama2-13B*	-	85 B/s
RWKV-1.5B*	-	20 KB/s
RWKV-7B*	-	4 KB/s
tszip (Bellard 2023)	-	180 KB/s
Tsf-200K	15 KB/s	48 KB/s
TsfXL-200K	29 KB/s	170 KB/s
RWKV-200K	750 KB/s	1.32 MB/s
L3TC-200K	1.30 MB/s	4.35 MB/s
L3TC-800K	980 KB/s	3.87 MB/s
L3TC-3.2M	633 KB/s	2.50 MB/s

Table 5: Decoding speed on different platforms.

tational resources, so their decoding speeds are not reported.

Table 5 shows that the Llama models operate at speeds measured in bytes per second (B/s), while the pretrained RWKV models are significantly faster, measured in kilobytes per second (KB/s). Our proposed methods *L3TC* are the fastest, reaching decoding speeds up to 1.30 MB/s on mobile devices. The speed gains of *L3TC-200K* over the default RWKV-200K are primarily due to our proposed outlier-aware tokenizer. This tokenizer merges frequent subwords, allowing more data to be processed in a single inference, whereas the default RWKV-200K uses character-based tokenizer and processes only one character at a time.

4.4 Ablation Study

Discussion on the Outlier-aware Tokenizer. We conduct experiments with different vocabulary sizes and coverage settings. As shown in the left sub-figure of Fig. 5 (the cover-

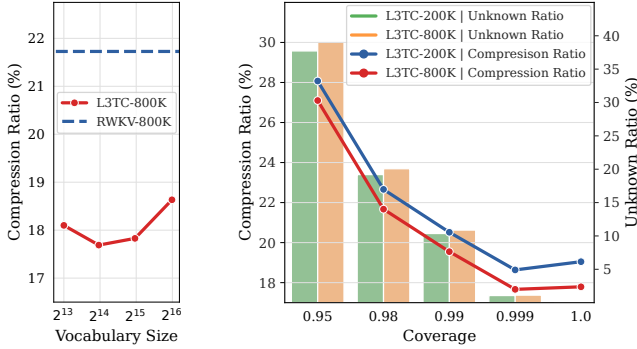


Figure 5: Discussion on the outlier-aware tokenizer.

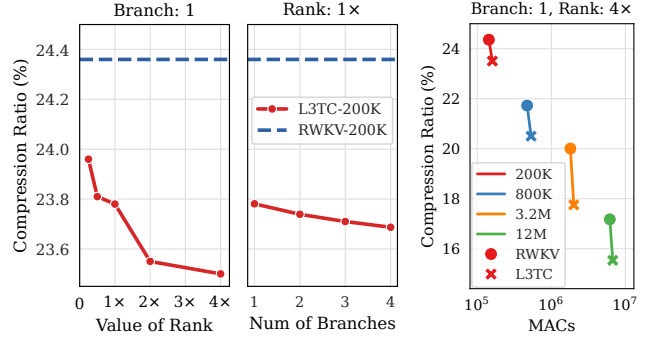
age for SPM-BPE is consistently set at 0.999), the compression ratios on enwik9 decrease as vocabulary size increases, adhering to the scaling law (Kaplan et al. 2020). However, excessively large vocabularies may require a larger model capacity, leading to a slight decline in compression performance when vocabulary sizes exceed 16K. Therefore, we empirically set the vocabulary size to 16K.

The right sub-figure of Fig. 5 demonstrates the relationship between coverage values and compression ratios using line graphs. The bar charts, whose y-axis is depicted on the right, indicate the proportion of unknown tokens’ size to the compressed data size (i.e., $\frac{8 \times N(x_i \notin \Gamma)}{R}$). As the coverage value increases, both unknown ratios and compression ratios exhibit a decreasing trend. However, the compression ratio reaches the lowest point when the coverage reaches 0.999, after which it slightly increases. This results validate the effectiveness of proposed outlier-aware tokenizer.

Discussion on High-rank Reparameterization. To examine the effect of rank r and the number of branches m , we conduct experiments as Fig. 6a. The left subfigure illustrates that given r ranging in $\{0.25, 0.5, 1, 2, 4\}$, compression ratios continuously decrease along with the increase of ranks. We use $r = 4$ in our experiments. The right subfigure reveals a minimal decrease in compression ratios with more branches, so we use only one branch in our experiments.

To further validate the effectiveness of high-rank reparameterization, we conduct experiments with default character-based tokenizers in Fig. 6b. The circles refer to the original unmodified RWKV models, while the crossed points represent the proposed L3TC models using HiRA with one branch and the rank of 4. Points of the same color correspond to models with equivalent sizes. Obviously, introducing HiRA significantly improves compression performance. It is noteworthy that reparameterization itself does not increase MACs; the marginal increase is due to the added linear shortcut parallel to the whole RWKV module.

Discussion on Decoding Speed. Batch size is a key factor to decoding speed but large batch size can not always improve the decoding speed proportionally, especially on resource-limited devices. We analysis the effect of batch size on iPhone12. As shown in Fig. 7, when batch size is rela-



(a) Effect of HiRA’s settings.

(b) Effect of HiRA.

Figure 6: Discussion on high-rank reparameterization.




Figure 7: Decoding speeds with different batch sizes.

tively small, the decoding speed is proportion to the batch size, but when the computational resources saturate (i.e. batch size is larger than 256), the decoding speed is not increased any more. Thus, we use a batch size of 256 in our settings when running on device.

5 Conclusion

In this paper, we propose a novel learned lossless low-complexity text compressor (L3TC) method. First, we conduct extensive experiments to compare various architectures and select RWKV as our backbone due to its fast decoding speed. Second, we propose an outlier-aware tokenizer that focuses on frequent tokens while bypassing infrequent outliers. Third, we propose a high-rank reparameterization strategy to enlarge the model capacity during training while not increasing inference complexity. Experiments demonstrate that our proposed L3TC achieves impressive lossless compression performance, with 48% size reductions compared to gzip. L3TC offers compression performance comparable to other learned compressors with $50\times$ model size reductions. Besides, L3TC is the fastest among all the learned compressors, with a decoding speed of up to MB/s.

Acknowledgments

This work was supported by the Ant Group Research Fund, the Fundamental Research Funds for the Central Universities, the National Natural Science Foundation of China (62102024, 62331014, 62431015), and the Shanghai Key Laboratory of Digital Media Processing and Transmissions, China.

References

- Apple. 2023. Guide to Core ML tools. <https://apple.github.io/coremltools/docs-guides/index.html#>. Accessed: 2024-08-10.
- Bellard, F. 2021. NNCP v2: Lossless data compression with transformer. *Technical report, Amarisoft*.
- Bellard, F. 2023. ts_zip: Text Compression using Large Language Models. https://bellard.org/ts_zip/. Accessed: 2024-08-10.
- Boutell, T. 1997. *PNG (Portable Network Graphics) Specification Version 1.0*. W3C.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J.; Le, Q. V.; and Salakhutdinov, R. 2019. Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context. Accessed: 2024-08-10, arXiv:1901.02860.
- Delétang, G.; Ruoss, A.; Duquenne, P.-A.; Catt, E.; Genewein, T.; Mattern, C.; Grau-Moya, J.; Wenliang, L. K.; Aitchison, M.; Orseau, L.; Hutter, M.; and Veness, J. 2024. Language Modeling Is Compression. In *Proceedings of International Conference on Learning Representations (ICLR)*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Ding, X.; Zhang, X.; Ma, N.; Han, J.; Ding, G.; and Sun, J. 2021. Reprvgg: Making vgg-style convnets great again. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13733–13742.
- Duda., J. 2013. Asymmetric numeral systems: entropy coding combining speed of Huffman coding with compression rate of arithmetic coding. ArXiv:1311.2540.
- Goldman, O.; Caciularu, A.; Eyal, M.; Cao, K.; Szpektor, I.; and Tsarfaty, R. 2024. Unpacking Tokenization: Evaluating Text Compression and its Correlation with Model Performance. Accessed: 2024-08-10, arXiv:2403.06265.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation*. *Neural Computation*.
- Hoffmann, J.; Borgeaud, S.; Mensch, A.; Buchatskaya, E.; Cai, T.; Rutherford, E.; Casas, D. d. L.; Hendricks, L. A.; Welbl, J.; Clark, A.; et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.
- Howard, P. G.; and Vitter, J. S. 1991. Analysis of arithmetic coding for data compression. In *Proceedings of Data Compression Conference (DCC)*.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. LoRA: Low-Rank Adaptation of Large Language Models. Accessed: 2024-08-10, arXiv:2106.09685.
- Huang, Y.; Zhang, J.; Shan, Z.; and He, J. 2024. Compression Represents Intelligence Linearly. Accessed: 2024-08-10, arXiv:2404.09937.
- Huffman, D. A. 1952. A method for the construction of minimum-redundancy codes. In *Proceedings of the IRE*.
- HutterPrize. 2006. 500'000€ Prize for Compressing Human Knowledge. <http://prize.hutter1.net/>. Accessed: 2024-08-10.
- Kaplan, J.; McCandlish, S.; Henighan, T.; Brown, T. B.; Chess, B.; Child, R.; Gray, S.; Radford, A.; Wu, J.; and Amodei, D. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.
- Knoll, B. 2023. Cmix version 20, a lossless data compression program. <http://www.byronknoll.com/cmixon.html>. Accessed: 2024-08-10.
- Lester, B.; Lee, J.; Alemi, A.; Pennington, J.; Roberts, A.; Sohl-Dickstein, J.; and Constant, N. 2024. Training LLMs over Neurally Compressed Text. Accessed: 2024-08-10, arXiv:2404.03626.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. Accessed: 2024-08-10, arXiv:1711.05101.
- MacKay., D. J. C., ed. 2003. *Information theory, inference, and learning algorithms*. Cambridge University Press.
- Mahoney, M. 2006a. PAQ8H: A High Compression Ratio Data Compression Program. <http://mattmahoney.net/dc/paq8h.zip>. Accessed: 2024-08-10.
- Mahoney, M. 2006b. Text8 Dataset. <http://mattmahoney.net/dc/textdata.html>. Accessed: 2024-08-09.
- Mahoney, M. 2024. Large Text Compression Benchmark. <https://www.mattmahoney.net/dc/text.html>. Accessed: 2024-08-10.
- Meta. 2015. ZSTD, Zstandard - Fast real-time compression algorithm. <https://github.com/facebook/zstd>. Accessed: 2024-08-10.
- Panayotov, V.; Chen, G.; Povey, D.; and Khudanpur, S. 2015. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 5206–5210. IEEE.
- Pasco., R. C. 1977. Source coding algorithms for fast data compression (ph.d. thesis abstr). *IEEE Trans. Inf. Theory*.
- Pasco., R. C. 1996. GZIP file format specification version 4.3. *RFC 1952*.
- Peng, B.; Alcaide, E.; Anthony, Q.; et al. 2023. RWKV: Reinventing RNNs for the Transformer Era. Accessed: 2024-08-10, arXiv:2305.13048.
- Schmidt, C. W.; Reddy, V.; Zhang, H.; Alameddine, A.; Uzan, O.; Pinter, Y.; and Tanner, C. 2024. Tokenization Is More Than Compression. Accessed: 2024-08-10, arXiv:2402.18376.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Neural Machine Translation of Rare Words with Subword Units. Accessed: 2024-08-10, arXiv:1508.07909.
- Seward, J. 2000. On the Performance of BWT Sorting Algorithms. *Proceedings of the IEEE Data Compression Conference 2000*.
- Shannon., C. E. 1948. A mathematical theory of communication. *Bell Syst. Tech. J.*
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; Rodriguez, A.; Joulin, A.; Grave, E.; and Lample, G. 2023a. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971*.

Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Valmeekam, C. S. K.; Narayanan, K.; Kalathil, D.; Chamberland, J.-F.; and Shakkottai, S. 2023. LLMZip: Lossless Text Compression using Large Language Models. Accessed: 2024-08-10, arXiv:2306.04050.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. Accessed: 2024-08-10, arXiv:1706.03762.

Xiph.Org Foundation. 2001. Free Lossless Audio Codec. <https://xiph.org/flac/>. Accessed: 2024-08-09.

Ziv, A., Jacob; Lempel. 1977. A Universal Algorithm for Sequential Data Compression. *IEEE Transactions on Information Theory*.