

---

# Language Models as Zero-shot Lossless Gradient Compressors: Towards General Neural Parameter Prior Models

---

Hui-Po Wang Mario Fritz  
CISPA Helmholtz Center for Information Security  
Germany  
`{hui.wang, fritz}@cispa.de`

## Abstract

Despite the widespread use of statistical prior models in various fields, such models for neural network gradients have long been overlooked. The inherent challenge stems from their high-dimensional structures and complex interdependencies, which complicate effective modeling. In this work, we demonstrate the potential of large language models (LLMs) to act as gradient priors in a zero-shot setting. We examine the property by considering lossless gradient compression – a critical application in distributed learning – that depends heavily on precise probability modeling. To achieve this, we introduce LM-GC, a novel method that integrates LLMs with arithmetic coding. Our technique converts plain gradients into text-like formats, enhancing token efficiency by up to 38 times compared to their plain representations. We ensure that this data conversion maintains a close alignment with the structure of plain gradients and the symbols commonly recognized by LLMs. Our experiments indicate that LM-GC surpasses existing state-of-the-art lossless compression methods, improving compression rates by 10% up to 17.2% across various datasets and architectures. Additionally, our approach shows promising compatibility with lossy compression techniques such as quantization and sparsification. These findings highlight the significant potential of LLMs as a model for effectively handling gradients. Code is available at <https://github.com/hui-po-wang/LM-GC>.

## 1 Introduction

Statistical prior models have been applied successfully in various fields, including image denoising and super-resolution [Ulyanov et al., 2018, Gandelsman et al., 2019], vision task adaptation [Wang et al., 2021], and low-resource language tasks [Baziotis et al., 2020, Brown et al., 2020]. However, their use in modeling neural network gradients has been largely neglected. The potential reasons for this oversight might include (1) the high-dimensional nature of gradients, which makes them less intuitive to analyze; (2) the difficulty of collecting representative gradient data; and (3) the significant challenge of ensuring generalizability to unseen data, given the substantial effort required.

Instead of developing a model from scratch, this work investigates the potential of leveraging pre-trained large-scale language models (LLMs) as gradient priors in a zero-shot setting. We explore this potential through the lens of lossless gradient compression, a vital application in federated and distributed learning environments. The success of this compression heavily depends on precise probability modeling. An effective statistical model can significantly improve compression efficiency, whereas an inaccurate model may lead to poorer compression outcomes and could even increase the data size post-compression.

To address this, we introduce LM-GC, an innovative coding scheme for gradient compression that integrates pre-trained large language models (LLMs) with arithmetic coding. Our method involves transforming gradients into text-like formats that are easier for LLMs to reason. Specifically, we convert the raw bit data of floating points into hexadecimal numbers and incorporate separators, such as spaces, to clarify the concept of floating points for LLMs. This serialized text is then processed by pre-trained tokenizers and LLMs to determine the probability of each token, which is subsequently utilized in arithmetic coding. Empirical evidence supports that these design choices significantly enhance gradient modeling and, consequently, compression efficiency.

Overall, our contributions are summarized below.

- We introduce LM-GC, a novel coding scheme that integrates LLMs with arithmetic coding. This method utilizes LLMs as powerful prior models for gradients, setting a new benchmark in state-of-the-art lossless gradient compression.
- LM-GC demonstrates that transforming raw gradients into formats that LLMs can understand significantly impacts their reasoning capabilities and token efficiency. Empirical evidence indicates that this approach can affect compression rates up to 70% with recognizable symbols and 40% with proper separators. These findings underscore the critical role of effective conversion in enhancing compression performance.
- Experimental results demonstrate that LM-GC significantly surpasses existing baselines, including PNG, FLAC, LZMA, GZIP, and FPZIP, by 10% to 17.2% across various architectures and datasets. Additionally, our approach complements existing lossy compression methods such as quantization and sparsification, paving the way for advanced gradient compression techniques.

## 2 Related work

**Large-scale language models.** Language models aim to model the relation between texts. This problem has been extensively studied in recent decades via various approaches such as statistical models [Jelinek, 1998] and recurrent neural networks [Hochreiter and Schmidhuber, 1997]. Recently, the emergence of transformer-based models [Vaswani et al., 2017] along with large-scale text corpora has revolutionized the entire field, driving research into large-scale language models (LLMs). Models, such as those from GPT [Achiam et al., 2023, Brown et al., 2020] and LLAMA [Touvron et al., 2023, Zhang et al., 2024, Geng and Liu, 2023] families, are capable of solving diverse tasks in natural languages and demonstrate incredible generalizability toward unseen novel tasks, even across modalities [Mirchandani et al., 2023, Gruver et al., 2024]. Notably, recent work by Deletang et al. [2024] also explores the use of language models as general compressors. Our goal is to investigate the potential of LLMs as a strong prior specifically for gradients. Additionally, we offer practical considerations for handling floating-point data when structures exist within the data to be compressed.

In this work, we demonstrate for the first time that LLMs can understand the structure of network gradients, accurately modeling their probability distribution in a fully zero-shot manner. We verify our finding by taking LLMs as priors for arithmetic coding, yielding state-of-the-art lossless gradient compression under various settings.

**Deep generative priors.** An ongoing research direction beyond traditional statistical modeling is learning a deep generative model from massive data and leveraging the model as a "deep" prior. The concept has been widely considered in many applications, such as image denoising and super-resolution [Ulyanov et al., 2018, Gandelsman et al., 2019], vision task adaptation [Wang et al., 2021, Chang et al., 2019], and low-resource language tasks [Baziotis et al., 2020, Brown et al., 2020]. Although strong priors can facilitate various downstream applications, training such models for gradients can be costly and challenging due to their high dimensionality. Additionally, the generalizability of these models is often a concern and may be limited to specific types of networks [Ha et al., 2016, Wang et al., 2024b]. Instead of training a model from scratch, our work explores the potential of using off-the-shelf LLMs as strong priors over gradients. This will minimize the cost of training deep prior models and may inspire applications like gradient denoising and anomaly detection.

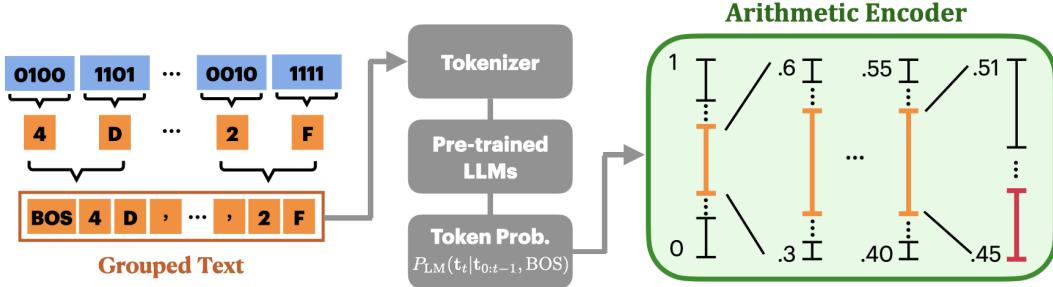


Figure 1: Overview of LM-GC. Our method initially converts every 4 bits into hexadecimal numbers and groups them with separators in between, e.g., commas in the figure. The grouped text is then input to a pre-trained, frozen tokenizer and LLM to produce the probability of each token. These probabilities are used for arithmetic encoding, where a line segment between 0 and 1 is repeatedly split according to the token probability until reaching a predefined maximum length. Any number from that region (e.g., the midpoint) can accurately represent the original data. We provide an example of how arithmetic coding works in Sec. 3.

**Gradient compression.** Gradient compression is a crucial technique, particularly in federated and distributed learning, where communication costs serve as the main bottleneck for scalability. Existing efforts have extensively studied lossy compression, which trades information precision for compression efficiency. For example, quantization [He et al., 2020, Alistarh et al., 2017, Bernstein et al., 2018] replaces floating points with fewer bits, while sparsification [Wangni et al., 2018, Alistarh et al., 2018] transmits only a subset of the original gradients. Other approaches explore novel optimization strategies such as progressive learning [Wang et al., 2022] and communicating synthetic images [Wang et al., 2024a, Xiong et al., 2023].

In contrast, lossless compression, which allows compressed data to be perfectly reconstructed without sacrificing information, is rarely investigated in the field of gradient compression nowadays. The challenge lies in developing a better statistical model for gradients. In this work, we demonstrate that LLMs can model the probability distribution of gradients in a zero-shot setting. Building on this finding, our method combines LLM-based modeling with arithmetic encoding and outperforms existing baselines such as PNG [Boutell, 1997], FLAC [Coalson, 2008], LZMA [Pavlov, 2019], GZIP [Deutsch, 1996], and FPZIP [Lindstrom and Isenburg, 2006], which are designed for modalities other than gradients. By integrating our approach with existing lossy compression techniques, we may pave the way for more advanced gradient compression schemes.

### 3 Background

In this work, we aim to explore the potential of using LLMs as prior for gradients and leverage lossless compression as an examination task. We review the essential background knowledge below.

**Lossless compression.** The fundamental principle of lossless compression is to reduce the size of data while ensuring it can be fully reconstructed. This is typically achieved by eliminating the statistical redundancy inherent in the data. Given a sequence of symbols  $s_{0:N} \in \mathcal{S}$  drawn from a probability distribution  $P_S$ , the objective is to devise a compression function  $g : \mathcal{S} \rightarrow \mathcal{C}$ . This function maps the original data  $s$  to a unique (decodable) binary code  $c$ , ensuring that the length of  $c$ , denoted by  $\ell(c)$ , is less than or equal to the length of  $s$ ,  $\ell(s)$ . The source coding theorem [Shannon, 2001] states that the expected minimum length of a coded message  $c$  cannot be shorter than the Shannon entropy of the original data, denoted as  $\ell(s) \geq H(\mathcal{S})$ . Here,  $H(\mathcal{S}) := \mathbb{E}_{s \sim P_S}[-\log_2 P_S(s)]$  represents the entropy. This implies that any compression resulting in a length shorter than  $H(\mathcal{S})$  necessarily involves loss of information, preventing perfect reconstruction of the original data.

**Arithmetic coding.** As a means to achieve lossless compression, arithmetic coding [Rissanen and Langdon, 1979] provides a nearly optimal message length  $H(\mathcal{S}) \leq \ell(c) \leq H(\mathcal{S}) + 2/\ell(s)$

on average [Sayood, 2017]. To implement it, one must employ a statistical probability model of the data, denoted  $P_{AC}$ . Ideally, this model  $P_{AC}$  should closely mirror the true distribution  $P_S$ . *The closer these distributions are, the more effective the compression performance will be.* Conversely, significant deviations can result in a compressed data length  $\ell(c)$  that exceeds the original data length  $\ell(s)$ . Notably, Most existing methods, such as CABAC, incorporate adaptive priors, meaning the probability  $P_{AC}$  adapts based on data context. However, as we will demonstrate in Sec. 5, these methods are not optimized for gradients and are thus outperformed by our zero-shot LLM prior.

The arithmetic encoder begins with an interval  $[0, 1)$  between 0 and 1. For each input symbol  $s$ , the interval is subdivided according to the probability  $P_{AC}(s)$ . The corresponding interval is then selected as the new interval. This process continues until the entire input stream is finished or reaches the maximum length. Any number existing in the final interval suffices to represent the compressed data. Similarly, the decoder takes the encoded output as input and can perfectly reconstruct the data by repeatedly looking up the intervals.

We provide an example illustrating how arithmetic coding works given a fixed statistical prior below.

**Example of arithmetic coding.** Consider a message consisting of only two symbols, A and B, where A occurs with probability  $P_{AC}(A) = 0.8$  and B with  $P_{AC}(B) = 0.2$ . The encoding interval gets subdivided into a larger part (0 to 0.8) for A and a smaller part (0.8 to 1) for B. If the message is "AAAB", the interval narrows from  $[0, 1]$  to  $[0, 0.8]$ , then  $[0, 0.64]$ , then  $[0, 0.512]$ , and finally  $[0.4096, 0.512]$  after processing the B. Any number (typically the midpoint for simplicity) within the final interval can represent the entire sequence. This number is then converted into a binary code, which is the compressed output. The final result is  $(0.4608)_{10} \rightarrow (0.01110101)_2$ , which takes only 8 bits compared to 4 bytes of storage for the ASCII format.

**Language models.** Language models are designed to model the relation between text symbols. Given a text stream,  $\mathcal{S} = \{s_i\}^N$ , consisting of N symbols, modern language models typically begin with a tokenization process  $f : \mathcal{S} \rightarrow \mathcal{T}$  (e.g., Byte-Pair Encoding [Sennrich et al., 2015]) that maps the entire stream to a set of K tokens  $\mathcal{T} = \{t\}^K$ . Then, the model predicts the probability as follows.

$$p(\mathbf{t}) = p(t_1, \dots, t_K) = \prod_{k=1}^K p(t_k | \text{BOS}, t_{<k}), \quad (1)$$

where the BOS token denotes a special token indicating the beginning of the sentence.

## 4 LM-GC

We introduce LM-GC, a method that integrates arithmetic coding with pre-trained large language models (LLMs) to address the lack of gradient-specific priors in arithmetic coding. It's important to note that LLMs are originally trained on extensive text corpora and do not encounter gradients or model parameters during this training. A significant challenge is enabling LLMs to comprehend the structure of gradients. Our method involves two main steps: serialization and compression. In serialization, we convert the 32-bit floating points of gradients into a format understandable by LLMs, which we call grouped text. This text is then fed into the LLMs, which predict the probability of each token in an autoregressive manner and thus accomplish compression using arithmetic coding.

**Serialization.** We first note that gradients are represented as 32-bit floating points, with values ranging from  $-3.40282347 \times 10^{38}$  to  $-1.17549435 \times 10^{-38}$ . Due to significant variations in their magnitudes and the often ambiguous importance of each gradient element, directly inputting these values into large language models (LLMs) is impractical. LLMs have a fixed token limit, and representing a single gradient in plain form would consume excessive tokens, compromising the context's depth.

To address this, our method, LM-GC, initially divides the floating points into several disjoint 4-bit partitions, which are then encoded into hexadecimal numbers as illustrated in Figure 1. This encoding strategy allows for a token savings of approximately 38 times compared to using plain gradients, particularly under extreme value conditions.

Furthermore, we organize every eight decoded hexadecimal numbers (equivalent to 4 bytes) by inserting a separator between them. This format provides LLMs with a structured representation of how a floating point number is typically presented. Our experiments demonstrate that separators are crucial in effectively modeling gradients, especially gradients derived from sophisticated architectures and datasets.

**Compression.** After serializing the gradients into grouped text  $\mathcal{S}$ , we process this text through a tokenizer to generate a set of tokens  $\mathcal{T}$ . These tokens are then fed into a pre-trained large language model (LLM), denoted by  $\mathcal{M}$ , which predicts the probability of each token as follows:

$$P_{\text{LM}}(\mathcal{T}) := \prod_{k=1}^K p(t_k | \text{BOS}, t_{<k}). \quad (2)$$

This equation indicates that the LLM sequentially predicts the probability of the next token, starting with the BOS (beginning of sequence) token, which is used to calculate the probability of the first token,  $P_{\text{LM}}(t_1) = \mathcal{M}(\text{BOS})$ . The BOS token serves primarily as a contextual cue and is not included in the compression.

During compression,  $P_{\text{LM}}(\mathcal{T})$  acts as the statistical model,  $P_{\text{AC}}$ , for arithmetic coding. In the decompression phase, the process begins with the BOS token, retrieving  $P_{\text{LM}}(t_1)$  to decode the first token. This decoding cycle continues until the maximum window size of the LLM is reached.

## 5 Experiments

### 5.1 Setup

**Datasets and models.** Our experiments consider three types of LLMs as the compressor, including Tinyllama 1.1B [Zhang et al., 2024], Openllama 3B [Geng and Liu, 2023], and LLAMA 2 7B [Touvron et al., 2023], ranging from a smaller to medium model size. All models can accept up to 4096 tokens. Unless stated otherwise, we use a context window size of 2048 by default. To ensure generalizability, we conduct experiments on four model architectures, including a three-layer convolution net (ConvNet), VGG-16 [Simonyan and Zisserman, 2015], ResNet [He et al., 2016], and vision transformer (ViT) [Dosovitskiy et al., 2021]. The models are trained on three datasets, MNIST [LeCun et al., 2010], CIFAR-10 [Krizhevsky et al., 2009], TinyImageNet [Le and Yang, 2015], under different settings. MNIST is a digit classification task containing 10 digits and 60000 images. On the other hand, CIFAR-10 and TinyImageNet are image classification tasks. CIFAR-10 contains 50000 images of 10 classes, while TinyImageNet contains 100000 images for 200 classes. All images are rescaled to 32 by 32 in the experiments.

**Evaluation protocols.** If a prior describes data well, lossless compression can achieve better efficiency. To test the efficiency of using LLMs as priors, we first train models on different datasets for 200 epochs, collecting gradients every 200 batch steps, resulting in a data pool of approximately 400 checkpoints for compression evaluation. Due to computational time constraints, we sub-sample 10 checkpoints from the pool for the subsequent experiments unless stated otherwise. All the experiments are repeated at least three times, and the standard deviations are reported accordingly. We measure compression efficiency by the compression rates defined as follows.

$$\text{Compression Rate (\%)} = 100 \times \frac{\text{Compressed Data Size}}{\text{Original Data Size}} \quad (3)$$

**Baselines.** We compare our method to state-of-the-art lossless compression techniques that originally targeted different data types. PNG [Boutell, 1997] is one of the most common lossless compression codecs for images. On the other hand, FLAC [Coalson, 2008] is a common audio compression format. Lastly, LZMA [Pavlov, 2019] and GZIP [Deutsch, 1996] are codecs used by 7-zip software and 7z compression format. FPZIP [Lindstrom and Isenburg, 2006] is proposed for scientific floating-point data, particularly suitable for data with up to 4D structures.

	Traditional codec		LM-GC (Ours)					
	Unchunked	Chunked	ISO	$H_n$	$H_s$	$H_c$	$H_{c+s}$	$H_{semi}$
PNG	43.30±1.3	49.18±1.1						
FLAC	52.37±0.6	50.46±0.6						
GZIP	42.42±0.3	47.10±0.4						
LZMA	41.91±0.0	47.36±0.1						
FPZIP	41.26±0.8	49.27±0.3						
Tinyllama 1.1B			117.38±0.0	36.30±0.8	38.83±0.4	38.40±0.6	38.46±0.1	43.45±0.6
Openllama 3B				71.85±0.2	37.07±0.1	32.32±0.3	34.31±0.6	33.07±0.5
LLAMA 2 7B				109.07±0.2	72.10±0.5	32.26±0.5	32.96±0.3	<b>32.21±0.8</b>
								32.78±0.4

Table 1: Gradient compression rate using PNG, FLAC, GZIP, LZMA, FPZIP, and our method with various language models. Our method considers different serializations including iso-8859-1 (ISO), hexadecimal numbers without separators ( $H_n$ ) and with spaces ( $H_s$ ), commas ( $H_c$ ), commas+spaces ( $H_{c+s}$ ), and semicolons ( $H_{semi}$ ) to group every four bytes from the same floating point.

**Implementation.** We implement our method in Pytorch and Huggingface. The checkpoints of pre-trained LLM models are loaded from the Huggingface hub. We adapted the arithmetic coding from Torchac to fit our application. We run our experiments on a cluster with NVIDIA A100 40GB GPUs and AMD EPYC 7402 24-Core Processor. All of the experiments can fit in one single A100.

## 5.2 Compression effectiveness

We first conduct compression experiments on gradients collected from a ConvNet trained on CIFAR-10 to show that LLMs can model gradients even without seeing such data during training. Our method considers three LLMs, namely Tinyllama, Openllama, and LLAMA 2, as the priors for arithmetic coding. We also consider 6 types of serialization, including decoding every byte with ISO-8859-1 (ISO), projecting every 4 bits to hexadecimal numbers without separators ( $H_n$ ), and with space ( $H_s$ ), commas ( $H_c$ ), space and commas ( $H_{c+s}$ ), and semicolons ( $H_{semi}$ ) as the separators. These settings outline the importance of serialization and its effect on gradient modeling. We report two settings for the baselines. The first is a *chunked* version, where the compressor sees a chunk of size 512 bytes every time, whereas the other one, namely the *unchunked* version, takes advantage of the pseudo infinitely large context length to yield the best statistical modeling.

Table 1 shows that our LM-GC consistently outperforms baseline codecs when serialization is properly managed. For example, ISO and  $H_n$  for LLAMA 2 perform worse than the baselines. In particular, ISO encodes gradients into symbols less familiar to LLMs, yielding up to 70% performance difference compared to settings like  $H_s$ . The lack of separators may confuse language models, causing performance degradation of 40% on LLAMA 2. These results highlight the crucial role of serialization in aiding LLMs' understanding. Furthermore, compression efficiency increases as the model size grows from 1.1B to 7B, suggesting that more sophisticated models may better understand the relationships between data elements, resulting in more effective compression.

## 5.3 Ablation study

In this section, we provide a series of ablation studies to provide insights into how design choices affect LLMs and the resulting prior models.

**Architectures.** To further understand the generalizability to different architectures and the effect of serialization, we continue with an experiment on different architectures. We extend the experiments to three additional architectures. VGG-16 contains deeper layers compared to ConvNets. ResNet-18 further introduces skip-connections and batch normalization, verifying our LM-GC on common design choices in modern machine learning. Lastly, ViT is built upon transformer blocks, showing that LLMs can reason beyond convolution layers. As shown in Table 2, we first observe that the performance of all methods drops as the models become more complicated, while our method remains the best among the baselines. This finding suggests that our method can better capture complex structures within the gradients. Moreover, we observe that serialization with separators generally

	Traditional codec					Ours (Tinylama 1.1B)			
	PNG	FLAC	GZIP	LZMA	FPZIP	$H_n$	$H_s$	$H_c$	$H_{c+s}$
ConvNet	43.30 $\pm$ 1.3	52.37 $\pm$ 0.6	42.42 $\pm$ 0.3	41.91 $\pm$ 0.0	41.26 $\pm$ 0.75	<b>36.30<math>\pm</math>0.8</b>	38.83 $\pm$ 0.4	38.40 $\pm$ 0.6	38.46 $\pm$ 0.1
VGG16	95.61 $\pm$ 0.2	-	91.91 $\pm$ 0.0	91.27 $\pm$ 0.1	89.15 $\pm$ 0.17	83.23 $\pm$ 0.0	<b>73.42<math>\pm</math>0.1</b>	75.32 $\pm$ 0.2	73.97 $\pm$ 0.1
ResNet18	97.22 $\pm$ 0.1	-	92.47 $\pm$ 0.0	91.72 $\pm$ 0.1	90.72 $\pm$ 0.07	83.20 $\pm$ 0.3	<b>73.57<math>\pm</math>0.1</b>	75.55 $\pm$ 0.3	73.95 $\pm$ 0.2
ViT	94.50 $\pm$ 0.4	-	89.20 $\pm$ 1.2	87.98 $\pm$ 1.2	89.77 $\pm$ 0.48	78.65 $\pm$ 3.3	<b>70.83<math>\pm</math>1.8</b>	72.60 $\pm$ 2.0	71.62 $\pm$ 1.7

Table 2: Gradient compression (%) for convolution neural networks (ConvNet), VGG-16, ResNet-18, and ViT trained on CIFAR-10.

	Traditional codec					LM-GC ( $H_s$ )	Impr.
	PNG	FLAC	GZIP	LZMA	FPZIP		
MNIST	50.05 $\pm$ 4.3	55.20 $\pm$ 1.7	45.05 $\pm$ 5.2	43.19 $\pm$ 1.3	44.62 $\pm$ 0.6	<b>39.38<math>\pm</math>1.4</b>	8.8%
CIFAR-10	43.30 $\pm$ 1.3	52.37 $\pm$ 0.6	42.42 $\pm$ 0.3	41.91 $\pm$ 0.0	41.26 $\pm$ 0.8	<b>38.83<math>\pm</math>0.4</b>	5.9%
TinyImageNet	96.08 $\pm$ 0.1	107.36 $\pm$ 0.0	92.18 $\pm$ 0.0	91.06 $\pm$ 0.1	86.88 $\pm$ 0.1	<b>71.90<math>\pm</math>0.0</b>	17.2%

Table 3: Compression effectiveness on MNIST, CIFAR-10, and TinyImageNet datasets. We use a Tinylama as the compressor to compress the gradients of ConvNets. The raw data are converted to hexadecimal numbers with spaces as the separator. The improvement (Impr.) over the best baseline highlights the capability of LM-GC in modeling complex gradients.

performs better than the one without separators. It outlines the importance of separators, especially when the data to be compressed becomes more intricate.

**Datasets.** The previous experiment suggests that LM-GC models gradients more accurately than the existing baselines, especially when considering complex structures. We further explore this dimension by considering two additional datasets, MNIST and TinyImageNet. Table 3 presents the result comparing our method with Tinylama to the baselines. Datasets like TinyImageNet introduce higher compression difficulty due to the complex task. However, LM-GC demonstrates consistently promising performance across all datasets. The improvement over the best baselines (FPZIP) increases as the dataset becomes sophisticated. This finding aligns with the result in Table 2 that our method is generalizable and better at capturing complex structures than the existing codecs that are not optimized for gradient compression.

**Context window size.** LM-GC takes LLMs as prior over gradients. One natural question is whether the LLMs really consider the context and yield accurate probability modeling. Ideally, similar to the traditional codec, if we provide a larger context window, the statistical model should be able to

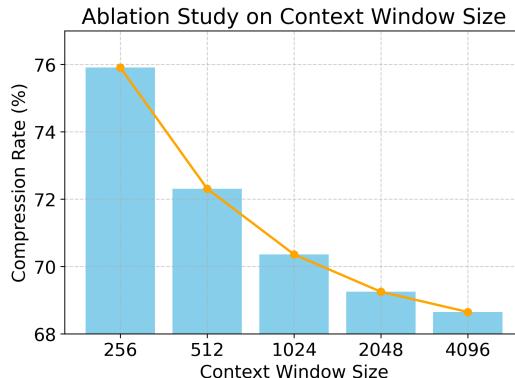


Figure 2: Compression rates of LLAMA 2-7B using context window sizes of 256, 512, 1024, 2048, and 4096. The compression rates improve as the context window increases.

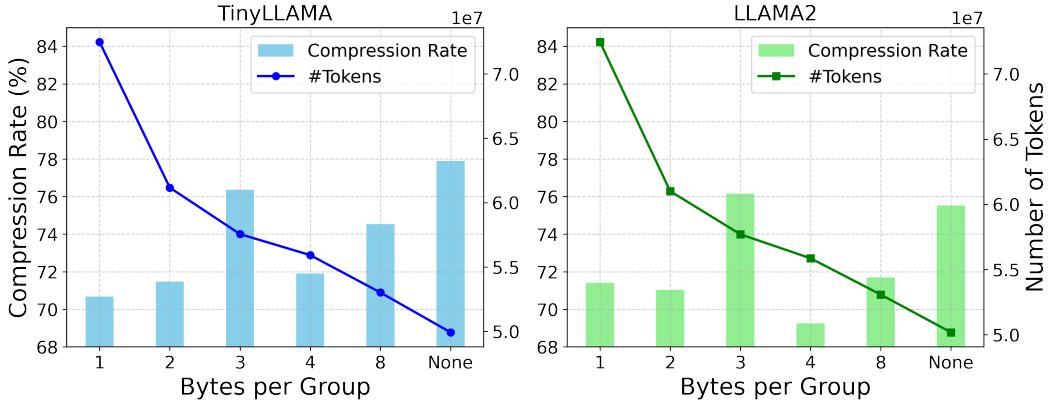


Figure 3: Ablation study on numbers of grouped bytes. We report the compression rates and the number of tokens yielded by different serializations. The settings that closely obey the data format perform better. However, smaller numbers yield higher computation overhead.

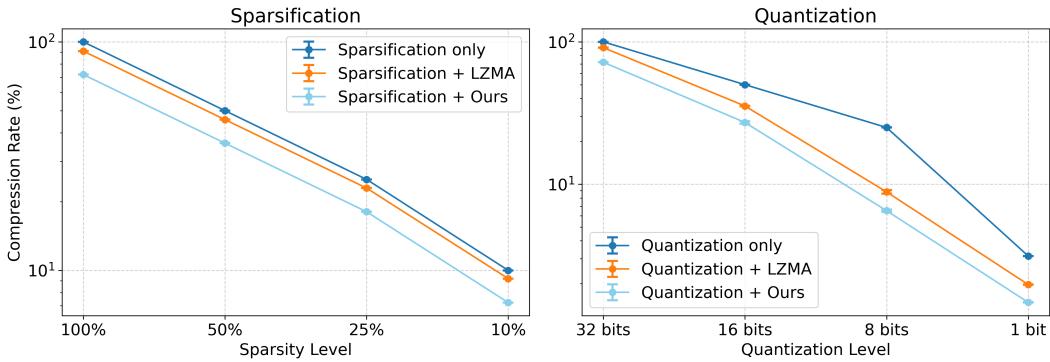


Figure 4: Compatibility analysis with sparsification (left) and quantization (right).

reason from the context and thus result in higher compression efficiency. Instead of using a default context window size of 2048 tokens, we conduct an ablation study in Fig. 2. The result shows that the performance drastically improves when the context window size increases, suggesting that LLMs indeed leverage the context. However, we note that the improvement seems to be saturated at the end. A larger context window also implies higher hardware resource demands, leaving a potential trade-off in practice.

**Byte grouping.** In addition to the decoding schemes analyzed in the previous experiments, we demonstrate that grouping converted text significantly affects performance. Recall that a floating point consists of 1 bit for the sign, 8 bits for the exponent, and 23 bits for the mantissa. Components with the same functionality should be grouped as closely as possible. To verify this hypothesis, we conducted experiments on TinyLLAMA and LLAMA 2 with bytes per group (BPG) set to 1, 2, 3, 4, 8, and none (i.e., no grouping, denoted as H<sub>n</sub>). The results in Fig. 3 show that BPG set to 1, 2, and 4 (our default setting) perform the best, while BPG equal to 3, which covers three components, and none perform worst. It indicates that serialization should resemble the structure of data to be compressed. Notably, although BPG equal to 1 and 2 performs well on both models, smaller BPG will add more separators and increase the total amount of tokens, introducing the additional computation overhead to the compression.

**Comparison to run-length encoding.** Lastly, we compare our method to run-length encoding, the simplest adaptive compression scheme, as shown in Table 4 in the appendix. The results indicate that although serialization may slightly improve compression rates, run-length encoding is ineffective

for compressing gradients. Combined with the earlier findings, this suggests that simple adaptive methods are unsuitable for handling complex yet structured gradient data.

#### 5.4 Compatibility

Gradient compression, a crucial technique in federated learning, typically involves lossy compression methods. We demonstrate that our LM-GC approach is compatible with lossy techniques such as quantization and sparsification. Specifically, we consider linear quantization, which uniformly discretizes value ranges according to the allotted bits. Given a vector  $v$  and  $n$  bits, the quantization process can be formalized as follows.

$$\bar{v} = \frac{v - \min v}{\max v - \min v} \times (2^n - 1). \quad (4)$$

In practice, only the indices  $I \in \{0, \dots, n\}$  for each element are communicated. Therefore, we map the data to the indices before conducting compression. Moreover, we consider sparsification, which selectively transmits a subset of gradients based on the specified proportion. When considering sparsification, it is important to note that the gradients remain as 32-bit floating points. For this experiment, we investigate quantization levels of 16, 8, and 1 bit (i.e., SignSGD [Bernstein et al., 2018]), and sparsification levels of 50%, 25%, and 10%.

We present a compatibility analysis in Fig. 4. The results indicate that integrating lossless compression techniques such as LZMA and LM-GC enhances compression rates beyond plain lossy compression. However, LZMA shows limited improvement across all settings, particularly with sparsification. In contrast, our method consistently delivers improvements across all settings, achieving notable compression rates in addition to lossy compression. These findings underscore the potential of LLMs as a prior for gradient compression, even with the incorporation of additional compression schemes, suggesting a promising new research direction in leveraging LLMs for compression.

### 6 Discussion and Limitation

**Throughput.** Despite the promising performance and generalizability, the throughput of LM-AC can be further optimized. Currently, our approach requires approximately 4 hours to compress just 28 MB. This bottleneck arises primarily from two components: LLMs and arithmetic coding. For LLMs, performance can be accelerated through techniques such as quantization [Frantar et al., 2023], faster attention mechanisms [Dao et al., 2022], KV cache Hooper et al. [2024], and model pruning Ma et al. [2023]. Looking ahead, one could explore distilling language models [Hsieh et al., 2023], as many functionalities may not be necessary during compression. Additionally, our implementation is significantly hindered by arithmetic coding and CPU limitations. Adopting a more efficient implementation, such as pure C++ programs, or utilizing CPUs with superior single-thread processing speeds could effectively mitigate these constraints.

**Broader impact.** Our work highlights the potential of leveraging pre-trained LLMs as priors for gradients. Immediately, this offers an advanced tool for gradient compression that reduces resource demands in federated and distributed learning environments. Over time, these priors could be utilized for gradient denoising, enhancing differential privacy training, or identifying adversarial gradients concealed within federated learning clients. However, this approach may also enable more subtle adversarial gradients, guided by these stronger priors.

### 7 Conclusion

We presented LM-GC, the first lossless gradient compressor that integrates arithmetic coding with LLMs as prior models for gradients. Our experiments show that pre-trained zero-shot LLMs are highly effective as gradient priors, setting a new state-of-the-art for gradient compression. Additionally, our findings indicate that the precise serialization of gradients substantially improves the reasoning abilities of LLMs and significantly impacts compression performance, warranting further exploration. The versatility of LM-GC sets the stage for developing more sophisticated gradient compression methods that directly incorporate LLMs. Overall, while our results in zero-shot settings are promising, the potential of expanding this approach to include few-shot learning, prompt engineering, and optimization of throughput efficiency remains open for further exploration.

## Acknowledgements

This work is partially funded by Medizininformatik-Plattform "Privatsphären-schützende Analytik in der Medizin" (PrivateAIM), grant No. 01ZZ2316G, and Bundesministeriums fur Bildung und Forschung (PriSyn), grant No. 16KISAO29K. The work is also supported by ELSA – European Lighthouse on Secure and Safe AI funded by the European Union under grant agreement No. 101070617. Moreover, The computation resources used in this work are supported by the Helmholtz Association's Initiative and Networking Fund on the HAICORE@FZJ partition. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the European Commission can be held responsible for them.

## References

J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

D. Alistarh, T. Hoefler, M. Johansson, N. Konstantinov, S. Khirirat, and C. Renggli. The convergence of sparsified gradient methods. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

C. Baziotis, B. Haddow, and A. Birch. Language model prior for low-resource neural machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.

J. Bernstein, Y.-X. Wang, K. Azizzadenesheli, and A. Anandkumar. signsgd: Compressed optimisation for non-convex problems. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.

T. Boutell. Png (portable network graphics) specification version 1.0. Technical report, 1997.

T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

W.-L. Chang, H.-P. Wang, W.-H. Peng, and W.-C. Chiu. All about structure: Adapting structural information across domains for boosting semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

J. Coalson. Free lossless audio codec, 2008. URL <https://xiph.org/flac>.

T. Dao, D. Fu, S. Ermon, A. Rudra, and C. Ré. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

G. Deletang, A. Ruoss, P.-A. Duquenne, E. Catt, T. Genewein, C. Mattern, J. Grau-Moya, L. K. Wenliang, M. Aitchison, L. Orseau, et al. Language modeling is compression. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024.

P. Deutsch. Gzip file format specification version 4.3. RFC, 1996.

A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2021.

E. Frantar, S. Ashkboos, T. Hoefler, and D. Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.

Y. Gandelsman, A. Shocher, and M. Irani. "double-dip": unsupervised image decomposition via coupled deep-image-priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

X. Geng and H. Liu. Openllama: An open reproduction of llama, May 2023. URL [https://github.com/openlm-research/open\\_llama](https://github.com/openlm-research/open_llama).

N. Gruver, M. Finzi, S. Qiu, and A. G. Wilson. Large language models are zero-shot time series forecasters. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

D. Ha, A. Dai, and Q. V. Le. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.

K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

Y. He, H.-P. Wang, M. Zenk, and M. Fritz. Cossqd: Communication-efficient federated learning with a simple cosine-based quantization. *arXiv preprint arXiv:2012.08241*, 2020.

S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 1997.

C. Hooper, S. Kim, H. Mohammadzadeh, M. W. Mahoney, Y. S. Shao, K. Keutzer, and A. Gholami. Kvquant: Towards 10 million context length llm inference with kv cache quantization. *arXiv preprint arXiv:2401.18079*, 2024.

C.-Y. Hsieh, C.-L. Li, C.-K. Yeh, H. Nakhost, Y. Fujii, A. Ratner, R. Krishna, C.-Y. Lee, and T. Pfister. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *arXiv preprint arXiv:2305.02301*, 2023.

F. Jelinek. *Statistical methods for speech recognition*. MIT press, 1998.

A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.

Y. Le and X. Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.

Y. LeCun, C. Cortes, and C. Burges. Mnist handwritten digit database. *ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist*, 2, 2010.

P. Lindstrom and M. Isenburg. Fast and efficient compression of floating-point data. *IEEE transactions on visualization and computer graphics*, 2006.

X. Ma, G. Fang, and X. Wang. Llm-pruner: On the structural pruning of large language models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

S. Mirchandani, F. Xia, P. Florence, B. Ichter, D. Driess, M. G. Arenas, K. Rao, D. Sadigh, and A. Zeng. Large language models as general pattern machines. In *Conference on Robot Learning*, 2023.

I. Pavlov. 7z format, 2019. URL <http://www.7-zip.org/7z.html>.

J. Rissanen and G. G. Langdon. Arithmetic coding. *IBM Journal of research and development*, 23(2): 149–162, 1979.

K. Sayood. *Introduction to data compression*. Morgan Kaufmann, 2017.

R. Sennrich, B. Haddow, and A. Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015.

C. E. Shannon. A mathematical theory of communication. *ACM SIGMOBILE mobile computing and communications review*, 5(1):3–55, 2001.

K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.

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

D. Ulyanov, A. Vedaldi, and V. Lempitsky. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

H.-P. Wang, N. Yu, and M. Fritz. Hijack-gan: Unintended-use of pretrained, black-box gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.

H.-P. Wang, S. Stich, Y. He, and M. Fritz. Progfd: effective, communication, and computation efficient federated learning by progressive training. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2022.

H.-P. Wang, D. Chen, R. Kerkouche, and M. Fritz. Fedlap-dp: Federated learning by sharing differentially private loss approximations. *Proceedings on Privacy Enhancing Technologies (PoPETs)*, 2024a.

K. Wang, Z. Xu, Y. Zhou, Z. Zang, T. Darrell, Z. Liu, and Y. You. Neural network diffusion. *arXiv preprint arXiv:2402.13144*, 2024b.

J. Wangni, J. Wang, J. Liu, and T. Zhang. Gradient sparsification for communication-efficient distributed optimization. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

Y. Xiong, R. Wang, M. Cheng, F. Yu, and C.-J. Hsieh. Feddm: Iterative distribution matching for communication-efficient federated learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

P. Zhang, G. Zeng, T. Wang, and W. Lu. Tinyllama: An open-source small language model. *arXiv preprint arXiv:2401.02385*, 2024.

## Appendix

### A Run Length Encoding

RLE (bits)	RLE ( $H_n$ )	RLE (ISO)	LM-GC ( $H_s$ )
450.28 $\pm$ 0.3	278.08 $\pm$ 0.2	198.57 $\pm$ 0.0	<b>71.90<math>\pm</math>0.0</b>

Table 4: Run length encoding results of gradients collected from ConvNets trained on TinyImageNet.

We additionally compare our method to run-length encoding (RLE). RLE compresses data by counting the consecutive symbols and replaces the original data with a series of (counts, symbol) tuples. It serves as a simple adaptive compression codec without knowing data characteristics. The experiment extends from Table 3, compressing gradients collected during training a ConvNet on TinyImageNet. We consider three types of dictionaries: binary, hexadecimal without separators ( $H_n$ , Table 1), and iso-8859-1 (extended ASCII to handle negative numbers). These methods use 1, 4, and 8 bits to represent symbols and always use 8 bits for counting. Note that this setting is favorable to RLE since gradient lengths can easily exceed 256 (8 bits).

The results are presented in Table 4. While different codebooks improve the efficacy of RLE, RLE failed to compress the data and even increase the data size. On the other hand, our method clearly outperforms RLE, indicating that simple adaptive priors are ineffective for gradients.

## NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer **[Yes]** , **[No]** , or **[NA]** .
- **[NA]** means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

**The checklist answers are an integral part of your paper submission.** They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "**[Yes]**" is generally preferable to "**[No]**", it is perfectly acceptable to answer "**[No]**" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "**[No]**" or "**[NA]**" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer **[Yes]** to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading “NeurIPS paper checklist”**,
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: **[Yes]**

Justification: As described in the abstract, we explore the potential of using LLMs as prior for gradients and take compression as an examination task.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: **[Yes]**

Justification: We have dedicated a section to discuss the potential limitations and solutions.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We briefly reviewed the background of lossless compression, but it's not our main contribution.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provided all implementation details and plan to release the source code upon publication.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: **[No]**

Justification: We will release the code upon publication.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: **[Yes]**

Justification: We have provided the required information in the main paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: **[Yes]**

Justification: We repeated all experiments at least three times to report the mean and standard deviation.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: **[Yes]**

Justification: Yes we have specified it.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: **[Yes]**

Justification: We have closely went through the code of ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: **[Yes]**

Justification: We dedicated one section to discuss it.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: **[NA]**

Justification: We do not release any models.

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We have cited all relevant works.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](http://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

## 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release any new assets at this point.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

## 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not include any related research mentioned in the guideline.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

**15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This work does not apply to IRB or any ethical reviews.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.