# **HBLLM: A Haar-Based Approach for Accurate Structured 1-Bit Quantized LLMs**

# Ningning Chen\* Sun Yat-sen University

Weicai Ye \* Sun Yat-sen University chennn27@mail2.sysu.edu.cn cai\_rcy@163.com

# Ying Jiang<sup>†</sup> Sun Yat-sen University jiangy32@mail.sysu.edu.cn

# **Abstract**

We introduce HBLLM, a wavelet-enhanced high-fidelity 1-bit post-training quantization method for Large Language Models (LLMs). By leveraging Haar wavelet transforms to enhance expressive capacity through frequency decomposition, HBLLM significantly improves quantization fidelity while maintaining minimal overhead. This approach features two innovative structure-aware grouping strategies: (1) frequency-aware multi-parameter intra-row grouping and (2)  $\ell_2$ -normbased saliency-driven column selection. For non-salient weights, a shared mean is employed across quantization groups within each frequency band to optimize storage efficiency. Experiments conducted on the OPT and LLaMA models demonstrate that HBLLM achieves state-of-the-art performance in 1-bit quantization, attaining a perplexity of 6.71 on LLaMA2-13B with an average weight storage of only 1.08 bits. Code available at: https://github.com/Yeyke/HBLLM.

# Introduction

In recent years, Large Language Models (LLMs) have achieved remarkable progress in natural language processing tasks. However, their massive parameter sizes—often reaching tens or even hundreds of billions—pose significant challenges for deployment on edge devices and in low-resource environments. To reduce the computational and memory burden of these models, a variety of compression techniques have been proposed, including quantization [12, 33, 35], pruning [11, 31], and knowledge distillation [19, 30]. Among them, Post-Training Quantization (PTQ) is widely adopted for its efficiency, requiring no additional training and having low deployment cost, especially in 1-bit quantization, which is considered a key approach for achieving extreme inference efficiency [13].




Figure 1: Average relative perplexity (normalized to FP16) on PTB, WikiText2, and C4 for LLaMA-1 family models, comparing LLM binarization methods and our HBLLM.

Although existing 1-bit PTQ methods [15, 17, 34] have achieved some success on base models such as GPT-2 and OPT, they tend to suffer from significant performance degradation—or even complete failure—when applied to more complex modern architectures like LLaMA3-8B [16]. To address this, recent studies have introduced several strategies to improve quantization fidelity:

<sup>\*</sup>Equal contribution.

<sup>&</sup>lt;sup>†</sup>Correspongding Author.

- **Group quantization**: divides the weight matrix into multiple groups for separate quantization. For instance, outlier-aware partitioning handles critical columns independently but can be constrained by partition design and scalability [15];
- **Residual approximation**: adds residual terms on top of primary quantization to partially recover errors [6], though this provides limited fidelity gains and introduces extra computation;
- Low-Rank Adaptation (e.g., LoRA): inserts low-rank modules to absorb quantization errors with some flexibility, like [34], but often shows sensitivity to hyperparameters;
- Global orthogonal transformations: apply global rotations in [1, 2, 5] before model compression to enhance representational capacity, but require expensive inverse transforms (e.g., matrix multiplications at  $\mathcal{O}(d^2)$  complexity for a d-dimension linear layer), leading to increased inference latency and energy consumption, making them impractical for deployment.

To overcome the structural trade-off between expressiveness and efficiency, we propose a novel 1-bit PTQ framework—**HBLLM**. This method is the first to integrate localized orthogonal transformations (i.e., Haar wavelets) into a BiLLM-style quantization process. Combined with structure-aware grouping, HBLLM significantly enhances expressive power under ultra-low bit budgets while maintaining negligible inverse transform cost and excellent compatibility with hardware-efficient inference.

Our main contributions are as follows:

- A **localized orthogonal transformation mechanism**: we apply a single Haar wavelet transform to decompose the weight matrix into high- and low-frequency components, improving binary expressiveness while reducing transform computation;
- Frequency-aware multi-parameter intra-row grouping: we introduce intra-row grouping in the frequency domain to capture structural patterns;
- $\ell_2$ -norm-based saliency-driven column selection: we propose an  $\ell_2$  norm-based ranking method to retain key columns using saliency metrics, effectively reducing quantization error;
- Intra-frequency-band mean sharing: for non-salient components, we introduce a mechanism that shares the mean across groups within the same row and wavelet band, reducing storage without sacrificing fidelity.

We conduct extensive experiments on OPT [37], LLaMA family [32] of LLMs. Results show that HBLLM achieves state-of-the-art performance under 1-bit quantization: Across language modeling tasks (C4, PTB, WikiText2), the perplexity ratio between HBLLM and the original FP16 model remains within the range of 1.2–2.2, shown in Fig 1, outperforming the next-best methods by 33%–66%; On 9 zero-shot QA benchmarks, HBLLM retains 73.8%–88.8% of the original model's accuracy; On modern architectures such as LLaMA3-8B, HBLLM remains stable with no performance collapse; Even with a lower average bit rate and memory usage than BiLLM and ARB-LLM<sub>RC</sub> [18], HBLLM outperforms both in overall task accuracy.

These results demonstrate that HBLLM significantly extends the applicability of 1-bit quantization, balancing extreme compression with high fidelity, and offers a new paradigm for deploying large-scale language models efficiently.

# 2 Related Work

# 2.1 1-Bit Post-Training Quantization

1-bit PTQ has emerged as a critical promising solution for deploying LLMs under extremely low bit budgets. Representative methods such as BiLLM [15] adopt a salient column separation mechanism, in which salient weights are quantized independently, while non-salient weights are grouped based on magnitude and quantized row-wise. ARB-LLM<sub>X</sub> [18] further introduces column-wise grouping and alternating refined binarization, achieving notable improvements in fidelity. Unlike [10, 34], BiLLM can accomplish PTQ tasks without intensive computation for knowledge distillation with multi-GPUs.

However, current methods face several key limitations: (1) They heavily rely on fixed thresholds or simple  $\ell_1$ -based heuristics for salient column selection, which are insufficient to capture sparse but significant activation outliers; (2) They fail to account for the structural asymmetry between row and

column dimensions in weight matrices, limiting their adaptability to complex model architectures; (3) They completely neglect frequency-domain information.

### 2.2 Evolution and Limitations of Grouping Strategies

To improve quantization flexibility and fidelity, some studies have proposed learnable or adaptive grouping strategies. For example, Mixture of Scales [17] introduces a Mixture-of-Experts (MoE) mechanism to assign scaling factor groups, and OneBitGPT [34] uses frequency masks to control quantization range sensitivity, and AWQ [3] identifies weights with the greatest impact on model predictions only based on activation outputs. However, these methods are generally effective only on unstructured tensors, rely on fine-grained distillation, and lack explicit frequency-domain awareness.

In addition, existing grouping strategies [15, 18] often apply uniform partitioning rules across the entire weight matrix, ignoring variations across different rows. This can lead to degraded expressiveness when quantizing models with significant inter-row diversity.

#### 2.3 Comparison Between Global Orthogonal Transforms and Local Wavelet Transforms

Orthogonal transforms have recently been adopted to improve LLM quantization. FrameQuant [1] and QuIP [5] utilize orthogonal transforms to enhance fidelity, but inference with such global transforms incurs high overhead, requiring  $\mathcal{O}(d^2)$  matrix multiplications [1] that cannot be fused into linear layers, leading to increased latency and energy cost.

By contrast, local orthogonal transforms such as the Haar wavelet [20] offer localized spectral sensitivity and have been widely applied in image compression, denoising, and edge detection [9, 14]. They can be efficiently implemented via lightweight local convolutions with negligible inference cost, making them well-suited for low-bit compression and edge deployment.

# 3 HBLLM: A Quantization Framework with Wavelet Transform and Frequency-Domain Grouping

# 3.1 Motivation and Core Challenges

Current mainstream 1-bit quantization methods face three key challenges in practice: (1) limited numerical expressiveness leading to high reconstruction error; (2) insufficient accuracy in salient column selection, failing to capture critical activation columns; (3) lack of structure-aware grouping strategies that adapt to heterogeneous model structures.

To characterize expressiveness under ultra-low bit settings, we introduce a new metric: the *cardinality* of the Inverse Quantization Set (CIQ), which measures the size of the discrete set of dequantized values within a row. CIQ serves as a unified indicator of how the above challenges constrain model fidelity. It acts both as a theoretical tool to analyze the limits of existing methods and as empirical evidence of the advantage of our proposed method.

Under 1-bit quantization, the CIQ of BiLLM and ARB-LLM $_{\rm X}$  is 8 and 10, respectively. When block size sets to 128, the CIQ upper bound of ARB-LLM $_{\rm X}$  can reach 128. In contrast, our method achieves a CIQ of up to 1024 after applying the Haar wavelet transform, significantly improving theoretical expressiveness. For more information on the benefits introduced by applying Haar transform, please refer to the appendix B and C.

Based on aboved analysis, we propose: (1) Haar wavelet transform to enhance expressive capacity by frequency decomposition; (2)  $\ell_2$ -norm-based saliency-driven column selection to prioritize critical columns; (3) frequency-aware multi-parameter intra-row grouping to capture structural patterns. We also introduce an intra-frequency-band mean sharing strategy and local convolution optimization to reduce storage and inference cost, thus forming a 1-bit PTQ framework **HBLLM** .

#### 3.2 Method Overview

We define the objective of HBLLM under the binary quantization setting for LLM weights. Specifically, the quantization targets the full-precision weight matrix  $\mathbf{W}_{FP} \in \mathbb{R}^{d \times d}$ , where a binary diagonal mask matrix  $\mathbf{M}_{sal} \in \{0,1\}^{d \times d}$  indicates which columns are selected as salient. The salient and




Figure 2: Overview of our HBLLM. The HBLLM quantization process consists of four steps: preparation, salient column selection, haar transform and quantization for the non-salient part, and quantization for the salient part. Since the salient columns are excluded from the Haar transform of the non-salient part, their positions must be filled before performing row-wise Haar transforms. This is handled by a process we refer to as FillAvg, where each missing column is filled with the average of its adjacent non-salient columns. For the non-salient part, HBLLM supports flexible choice between row-wise (HBLLM-row) and column-wise (HBLLM-col) transforms. The salient part undergoes column-wise Haar transformation followed by HaarQuant for quantization.

non-salient parts are quantized in the Haar domain, and their respective quantized Haar coefficients are denoted by  $\widehat{\mathbf{W}}_{B}^{sal}$  and  $\widehat{\mathbf{W}}_{B}^{non-sal}$ . These are then reconstructed using inverse Haar transforms  $\mathcal{H}_{1}^{-1}$  and  $\mathcal{H}_{2}^{-1}$ .

The reconstruction objective of HBLLM is twofold. For the quantization of a matrix layer  $\mathbf{W}$ , the objective expressed in the Frobenius norm is formulated as:

$$\min_{\widehat{\mathbf{W}}} \left\| \mathbf{W} \mathbf{X} - \widehat{\mathbf{W}} \mathbf{X} \right\|_{F}^{2}, \tag{1}$$

where X is the input of the matrix layer. For quantization of a matrix block  $W_{FP}$  of W, the object is:

$$\min_{\mathbf{M}_{\text{sal}},\widehat{\mathbf{W}}_{\text{B}}^{\text{non-sal}}} \left\| \mathbf{W}_{\text{FP}} - \mathbf{M}_{\text{sal}} \mathcal{H}_{1}^{-1} \left( \widehat{\mathbf{W}}_{\text{B}}^{\text{sal}} \right) - (\mathbf{I} - \mathbf{M}_{\text{sal}}) \mathcal{H}_{2}^{-1} \left( \widehat{\mathbf{W}}_{\text{B}}^{\text{non-sal}} \right) \right\|_{F}^{2}. \tag{2}$$

When  $\mathcal{H}_1 = \mathcal{H}_2$  are fixed Haar transforms, this formulation simplifies to a quantization problem entirely in the Haar domain. In this case, the objective is the same to that of BiLLM. Layer-level quantization is commonly tackled with the GPTQ algorithm [12].

We emphasize that our approach does not aim to solve this objective function via explicit optimization. Instead, this formulation serves as a conceptual framework that guides our method design. The actual quantization process is based on a set of heuristics and structure-aware strategies that approximate this objective in a computationally efficient and scalable manner.

**Quantization Pipeline Overview.** HBLLM integrates the Haar transform into a BiLLM-style quantization pipeline (see Algorithm 1 and Figure 2), consisting of the following key steps:

 Preparation Phase: Compute the column-wise importance scores using a Hessian-based saliency metric.

# 2. Salient Column Selection and Quantization(SALIENT):

- Sort columns by their  $\ell_2$  norm significance.
- Select top-K salient columns and determine  $\mathbf{M}_{\text{sal}}$ .
- $\widehat{\mathbf{W}}_{R}^{sal} = \text{HaarQuant}(\mathbf{M}_{sal}\mathbf{W}).$
- Choose the subset with the lowest quantization error.

# 3. Non-Salient Region Quantization:

- Fill the missing values in salient columns using adjacent averages (FillAvg).
- $\widehat{\mathbf{W}}_{B}^{\text{non-sal}} = \text{HaarQuant}\left(\mathbf{M}_{\text{sal}}\mathbf{W}_{\text{filled}} + (\mathbf{I} \mathbf{M}_{\text{sal}})\mathbf{W}\right)$ , where  $\mathbf{W}_{\text{filled}}$  is from FillAvg.

# 4. Adjustment and Refinement:

• 
$$\widetilde{\mathbf{W}} = \mathbf{M}_{\text{sal}} \left( \mathbf{W} - \mathcal{H}^{-1} \left( \widehat{\mathbf{W}}_{B}^{\text{non-sal}} \right) \right)$$
.

• 
$$\widehat{\mathbf{W}}_{B}^{sal} = \text{HaarQuant}\left(\widetilde{\mathbf{W}}\right)$$
.

Algorithm 1 Framework of HBLLM: Details of each function are shown in Algorithm E.1

```
func Row-HaarQuant(\mathbf{W}, rows\{\cdot\})
func HBLLM(\mathbf{W}, \mathbf{X}, \beta, \lambda)
Input: \mathbf{W} \in \mathbb{R}^{n \times m} - weight matrix
                                                                                                                1: \mathbf{W}_{\text{filled}} \leftarrow \text{FillAvg}(\mathbf{W}_{:,j \notin \text{rows}}, \text{rows}\{\cdot\})
              \mathbf{X} \in \mathbb{R}^{r \times d} - calibration data
                                                                                                                2: \mathbf{B}_{\text{filled}} \leftarrow \text{HaarQuant}(\mathbf{W}_{\text{filled}}, \text{ROW})
              \beta - block size
                                                                                                                3: \widehat{\mathbf{W}} \leftarrow \mathbf{W} - \mathbf{B}_{\text{filled}}
              \lambda - hessian regularizer
                                                                                                                4: \mathbf{B}_{\text{salient}} \leftarrow \text{HaarQuant}(\widehat{\mathbf{W}}_{:,i \in \text{rows}}, \text{COL})
Output: B - haared binarized weights
                                                                                                                5: \mathbf{B} \leftarrow \mathbf{B}_{\text{salient}} + \mathbf{B}_{\text{filled}}
  1: \mathbf{H} \leftarrow 2\mathbf{X}\mathbf{X}^{\top} // \ell^2 error hessian matrix
                                                                                                                6: return B
 2: \mathbf{H}^c \leftarrow \text{Cholesky}((\mathbf{H} + \lambda \mathbf{I})^{-1})
                                                                                                             func Col-HaarQuant(\mathbf{W}, rows\{\cdot\})
 3: \mathbf{B} \leftarrow \mathbf{0}_{n \times m}
                                                                                                                1: \mathbf{B}_{unsalient} \leftarrow HaarQuant(\mathbf{W}_{:,j \notin rows}, COL)
 4: for b = 0, \beta, 2\beta, ..., N do
                                                                                                                2: \widehat{\mathbf{W}} \leftarrow \mathbf{W} - \mathbf{B}_{\text{filled}}
               \mathbf{W}^b \leftarrow \mathbf{W}_{:,b:b+\beta}
 5:
                                                                                                                3: \mathbf{B}_{\text{salient}} \leftarrow \text{HaarQuant}(\widehat{\mathbf{W}}_{:,j \in \text{rows}}, \text{COL})
               rows{\cdot} \leftarrow SALIENT(\mathbf{W}_{:,b:b+\beta}, \mathbf{H}^c)
 6:
                                                                                                                4: B ← B<sub>salient</sub> + B<sub>unsalient</sub>
               if Row-HBLLM then
 7:
                       \mathbf{B}_{:,b:b+\beta} \leftarrow \text{Row-HaarQuant}(\mathbf{W}^b, \text{rows}\{\cdot\})^{5:} return \mathbf{B}
 8:
               else if Col-HBLLM then
 9:
                       \mathbf{B}_{:,b:b+\beta} \leftarrow \text{Col-HaarQuant}(\mathbf{W}^b, \text{rows}\{\cdot\})
10:
               \mathbf{E} \leftarrow (\mathbf{W}_{:,b:b+\beta} - \mathbf{B}_{:,b:b+\beta}) / \mathbf{H}_{b:b+\beta,b:b+\beta}^c
11:
               \mathbf{W}_{:,b+\beta:} \leftarrow \mathbf{W}_{:,b+\beta:} - \mathbf{E} \cdot \mathbf{H}_{b:b+\beta:b+\beta:}^c
12:
13: return B
```

# 3.3 HaarQuant: One-Bit Quantization in the Wavelet Domain

To boost expressiveness, we apply Haar wavelet transform to the weight matrix of linear layers, generating a frequency-domain coefficient matrix, followed by group-wise 1-bit quantization.

To address limited numerical expressiveness, HBLLM introduces the HaarQuant algorithm. Haar-Quant consists of three stages.

**Haar Transform.** A row of weights W is decomposed into low- and high-frequency coefficients via 1D Haar transform  $\mathcal{H}$ :

$$\widehat{\mathbf{W}} = \mathcal{H}(\mathbf{W}) = \left[\mathcal{H}_{\text{low-pass}}(\mathbf{W}), \mathcal{H}_{\text{high-pass}}(\mathbf{W})\right], \tag{3}$$

where  $\widehat{\mathbf{W}}$  is the Haar coefficient of  $\mathbf{W}$ ,  $\mathcal{H}_{low-pass}\left(\mathbf{W}\right)$  and  $\mathcal{H}_{high-pass}\left(\mathbf{W}\right)$  are low- and high-frequency coefficients, respectively.

**Frequency-Aware Multi-Parameter Intra-Row Grouping.** For each row, boundary candidates determined by the row are enumerated, and the best grouping with minimal quantization error is selected. Furthermore, we split the rows by frequency bands. This adaptive strategy captures intra-row structural differences better than global uniform boundaries used in BiLLM.

**Coefficient Quantization.** Each group  $\widehat{\mathbf{W}}_{\mathrm{FP}}$  is quantized using sign-based binarization centered on its mean:

$$\widehat{\mathbf{W}}_{\mathrm{B}} = \alpha \cdot \mathrm{sign}(\widehat{\mathbf{W}}_{\mathrm{FP}} - \mu),\tag{4}$$

where  $\alpha \in \mathbb{R}^d$  is the row-wise scaling factor and  $\mu$  is the group-wise mean and  $\widehat{\mathbf{W}}_{\mathrm{B}}$  is the result.

# 3.4 Structure-aware Grouping Strategies

To enhance the fidelity and adaptability of binary quantization under structural constraints, HBLLM introduces two structure-aware grouping strategies that operate along both column and row dimensions of the weight matrix.

**Saliency-Driven Column Selection via**  $\ell_2$  **Norm.** This strategy is used during salient column identification and quantization to overcome the limitations of prior heuristics based on fixed thresholds or simple magnitude criteria.

- Columns are ranked by their ℓ₂-norm scores, which correlate with their overall contribution to activation magnitude.
- The top-K columns are selected as salient and quantized in the Haar-transformed domain using column-wise transforms.

This approach helps preserve activation-critical directions, especially those dominated by outlier weights.

**Frequency-Aware Multi-Parameter Intra-Row Grouping.** This strategy is used during Haar domain quantization, where conventional row grouping lacks sensitivity to structural variations in weight distributions.

- Each row is first decomposed into high- and low-frequency components based on Haar subbands.
- Within each frequency band, coefficients are adaptively split into dense and sparse groups using band-specific, data-driven thresholds.

This grouping effectively doubles the number of quantization subgroups per row, enabling finer granularity and better error control.

Together, these strategies facilitate fine-grained, structure-preserving quantization across both dimensions of the weight matrix. To further guide saliency-based partitioning, we adopt the parameter importance metric used in BiLLM, defined as:  $s_i = w_i^2/[\mathbf{H}^{-1}]_{ii}^2$ , where  $\mathbf{H}$  denotes the Hessian matrix of the layer,  $w_i$  is the full-precision value of the i-th parameter, and  $[\mathbf{H}^{-1}]_{ii}$  is the i-th diagonal entry of the inverse Hessian.

This metric reflects the relative sensitivity of the loss to changes in each parameter: higher values indicate greater influence on the model's output, and thus prioritize that weight for accurate reconstruction.

#### 3.5 Intra-frequency-band Mean Sharing

To reduce storage overhead, HBLLM shares a single mean value among 2 groups in the same frequency band within each row:  $\mu_{\text{shared}} = \frac{1}{n_1 + n_2} \left( \sum_{i=1}^{n_1} x_i + \sum_{j=1}^{n_2} y_j \right)$ . It not only reduces perparameter storage by 0.25 bits, but also maintains accuracy even slightly improving downstream task performance. This optimization achieves a trade-off between compression and accuracy, improving deployment viability.

# 3.6 Efficient Haar Implementation via Local Convolutions

Instead of costly matrix multiplication, HBLLM implements Haar transform using fixed local convolutions. There are only two predefined 1D kernels, [1/2,1/2] and [1/2,-1/2], whose kernel size is 2. Furthermore, it can be hardcoded into the model for zero runtime initialization and no training or storage is needed. In complexity comparison, HBLLM needs  $\mathcal{O}(d)$  operations via convolutional sliding window, while FrameQuant needs  $\mathcal{O}(d^2)$  operations. As a result, HBLLM significantly lowers inference cost and is ideal for edge deployment.

Table 1: Comparison of perplexity and average accuracy across models and methods

	LLaMA1		D	omlovit	. I		· <u> </u>	LLaMA3		D,	mlovita	ı, I	
		***		erplexity		<b>AvgQA</b> ↑			*** * **		rplexity		<b>AvgQA</b> ↑
Size	Method	W-bits	C4	Wiki2	PTB		Size	Method	W-bits	C4	Wiki2	PTB	
	FullPrecision	16.00	6.71	5.68	35.80	65.62		FullPrecision	16.00	11.90	8.29	13.07	68.94
	FrameQuant	2.20	10.89	9.96	104.7	56.19		FrameQuant	2.20	28.44	23.36	40.33	52.27
7B	PB-LLM	1.70	90.19	113.4	830.0	35.71	OD.	PB-LLM	1.70	111.7	141.5	171.1	36.83
/ <b>D</b>	BiLLM	1.09	43.74	44.85	369.3	40.01	8B	BiLLM	1.06	53.67	56.24	81.27	41.84
	ARB-LLM <sub>X</sub>	1.09 1.09	22.80 15.13	24.70 13.45	240.5 155.8	45.65 52.23		ARB-LLM <sub>X</sub>	1.06	48.45 34.44	37.90 30.24	52.59 45.23	43.40 49.08
	ARB-LLM <sub>RC</sub> HBLLM-row	1.09	9.49	8.82	88.86	57.48		ARB-LLM <sub>RC</sub> HBLLM-row	1.06 1.06	20.09	16.18	<b>22.83</b>	54.80
	HBLLM-col	1.00	10.38	9.67	117.7	54.03		HBLLM-col	1.00	22.18	17.80	26.38	51.43
	FullPrecision	16.00	6.24	5.09	25.36	68.09		FullPrecision	16.00	6.61	2.85	7.74	74.62
	FrameQuant	2.20	8.79	7.84	50.69	60.69		FrameQuant	2.20	N/A	N/A	N/A	N/A
13B	PB-LLM	1.70	38.41	46.02	190.2	40.39	70B	PB-LLM	1.70	33.56	28.93	44.38 129.5	47.45
130	$BiLLM$ $ARB-LLM_X$	1.10 1.10	13.93 N/A	14.99 N/A	69.75 N/A	50.89 N/A	70 <b>B</b>	$BiLLM$ $ARB-LLM_X$	1.09 1.09	385.8 N/A	137.6 N/A	129.5 N/A	34.18 N/A
	ARB-LLM <sub>RC</sub>	1.10	10.68	10.19	43.85	59.58		ARB-LLM <sub>RC</sub>	1.09	12.80	10.24	12.76	63.90
	HBLLM-row	1.09	7.62	6.68	34.94	62.57		HBLLM-row	1.08	10.87	8.08	11.44	56.45
	HBLLM-col	1.00	7.77	6.98	37.62	61.25		HBLLM-col	1.00	13.69	9.09	14.26	55.89
	FullPrecision	16.00	5.62	4.10	21.35	71.06		ОРТ		P	erplexit	y↓	AvgQA↑
	FrameQuant PB-LLM	2.20 1.70	7.35 21.73	6.32 25.87	28.69 127.1	65.13 47.22	Size	Method	W-bits	C4	Wiki2	PTB	841
30B	BiLLM	1.11	10.27	10.55	41.76	58.07	-	FullPrecision	16.00	13.45	14.62	16.41	52.54
	ARB-LLM <sub>X</sub>	1.11	N/A	N/A	N/A	N/A		FrameQuant	2.20	24.29	27.15	30.45	44.48
	ARB-LLM <sub>RC</sub>	1.11	8.49	7.79	30.98	64.49		PB-LLM	1.70	186.9	309.0	286.3	33.44
	HBLLM-row HBLLM-col	1.10	6.88	5.82	25.95	<b>66.76</b>	1.3B	BiLLM	1.09	56.24	68.43	119.2	38.39
	HDLLM-COI	1.00	7.03	6.03	26.65	64.86		$ARB-LLM_X$	1.09	43.23	53.55	67.96	41.42
	FullPrecision	16.00	5.31	3.53	21.11	72.27		ARB-LLM <sub>RC</sub>	1.09	24.23	28.77	33.32	45.28
	FrameQuant	2.20	6.69	5.55	27.48	68.58		HBLLM-row	1.07	19.30	21.68	25.34	46.35
	PB-LLM	1.70	12.66	12.76	99.67	62.48		HBLLM-col	1.00	21.92	24.08	27.28	44.70
65B	BiLLM	1.10	9.26	8.58	41.93	62.05		FullPrecision	16.00	12.06	12.47	14.61	54.95
	ARB-LLM <sub>X</sub>	1.10	N/A	N/A	N/A	N/A		E	2.20	117.06	10.24	22.60	1 40 50
	$ARB-LLM_{RC}$	1.10	7.48	6.47	29.14	68.53		FrameQuant PB-LLM	2.20 1.70	17.86 165.1	18.24 216.8	22.60 160.4	<b>49.58</b> 37.62
	HBLLM-row	1.09	6.28	5.07	24.11	69.18	2.7B	BiLLM	1.10	42.92	55.75	103.2	40.02
	HBLLM-col	1.00	6.44	5.26	30.38	67.83		ARB-LLM <sub>X</sub>	1.10	30.02	34.15	41.35	44.60
	LLaMA2		P	erplexity	y↓	AviaOAA	-	ARB-LLM <sub>RC</sub>	1.10	18.02	19.53	24.46	49.53
Size	Method	W-bits	C4	Wiki2	PTB	AvgQA↑		HBLLM-row HBLLM-col	1.09 1.00	15.70 17.28	<b>16.85</b> 18.80	<b>19.54</b> 22.63	48.80 48.56
	FullPrecision	16.00	8.66	6.94	37.86	65.54		FullPrecision	16.00	10.68	10.86	12.73	58.95
	FrameQuant	2.20	14.66	13.34	177.1	52.75		FrameQuant	2.20	14.53	14.59	18.71	53.77
	PB-LLM	1.70	63.95	55.40	486.2	36.54		PB-LLM	1.70	122.9	206.7	222.3	34.87
7B	BiLLM	1.08	33.97	31.38	373.0	42.11	6.7B	BiLLM	1.11	39.96	54.91	90.10	37.40
	$ARB-LLM_X$	1.08	26.55	21.74	314.2	45.41		$ARB-LLM_X$	1.11	19.39	19.50	24.78	49.79
	ARB-LLM <sub>RC</sub>	1.08	17.87	15.85	462.2	46.71		ARB-LLM <sub>RC</sub>	1.11	14.29	15.16	17.92	53.76
	HBLLM-row	1.07	11.75	10.52	89.23	57.74		HBLLM-row	1.10	12.56	13.04	15.26	56.17
	HBLLM-col	1.00	12.51	11.33	150.6	54.09		HBLLM-col	1.00	13.29	13.67	15.70	54.44
	FullPrecision	16.00	6.18	4.88	43.02	69.18		FullPrecision	16.00	10.16	10.13	11.89	58.41
	FrameQuant	2.20		7.80	109.3			FrameQuant		12.26	12.51	14.59	
13B	PB-LLM BiLLM	1.70 1.08	313.4 22.17	289.4 19.57	934.4 303.4	32.91 46.76	13B	PB-LLM BiLLM	1.70 1.13	42.89 17.01	18.34	94.98 21.56	39.50 49.82
130	ARB-LLM	1.08	N/A	N/A	N/A	N/A	131	ARB-LLM <sub>X</sub>	1.13	N/A	N/A	N/A	N/A
	ARB-LLM <sub>RC</sub>	1.08	11.90	10.98	151.8	57.35		ARB-LLM <sub>RC</sub>	1.13	12.60	13.14	15.14	55.35
	HBLLM-row	1.07	7.82	6.71	61.75	63.61		HBLLM-row	1.12	11.47	11.72	13.78	55.91
	HBLLM-col	1.00	8.28	7.00	69.74	62.04		HBLLM-col	1.00	11.71	12.34	14.13	55.66
	FullPrecision	16.00	5.24	3.32	21.49	72.96		FullPrecision	16.00	9.60	9.56	11.50	62.09
	FrameQuant	2.20	N/A	N/A	N/A	N/A		FrameQuant	2.20	10.92	11.15	13.25	59.62
705	PB-LLM	1.70	N/A	N/A	N/A	54.26	205	PB-LLM	1.70	21.60	28.62	45.63	46.14
70B	BiLLM	1.09	15.57	15.86	71.03	55.81	30B	BiLLM	1.06	13.43	13.44	16.66	54.22
	ARB-LLM <sub>X</sub>	1.09	N/A	N/A	N/A	N/A		ARB-LLM <sub>X</sub>	1.06	N/A	N/A	N/A	N/A
	ARB-LLM <sub>RC</sub>	1.09	7.26	6.00	28.43	68.77		ARB-LLM <sub>RC</sub> HBLLM-row	1.06	11.18	10.94	13.27	58.59
	HBLLM-row HBLLM-col	1.08 1.00	<b>6.18</b> 6.63	<b>4.82</b> 5.04	<b>24.69</b> 26.31	<b>70.01</b> 68.61		HBLLM-row HBLLM-col	1.06 1.00	10.41 10.53	<b>10.13</b> 10.29	<b>12.58</b> 12.75	<b>60.04</b> 58.91
	TDEEM-COI	1.00	0.05	2.04	20.51	00.01		TIDEENI-COI	1.00	10.55	10.27	12.73	C 1 0 0

*Note*: All methods are calibrated on C4 with 128 samples and a sequence length of 2048. A block size of 128 is used for channel-wise quantization, as commonly done in prior work. N/A: ARB-LLM<sub>X</sub> method cannot run on a single 3090 GPU - 24GB. W-bits is the average weight overhead per weight. For more details, please refer to the appendix D.

# 4 Experiments

# 4.1 Experimental Settings

Models and Evaluation Datasets. In our study, we evaluate HBLLM on various models, including those from the OPT, LLaMA-1, LLaMA-2, and LLaMA-3, as well as the recently introduced f-R1-Distill-Llama-8B. Specifically, we utilize the OPT models with 1.3B and 2.7B parameters, the LLaMA-1 and LLaMA-2 models with 7B and 13B parameters for our evaluations, and the LLaMA-3 model with 8B parameters. We measure language modeling capabilities of these models by evaluating their perplexity on the C4[26], WikiText2[22] and PTB[21] datasets. Additionally, we assess zero-shot accuracy on various Common Sense Reasoning Tasks such as PIQA[4], BoolQ[7], OpenBookQA[23], WinoGrande[28], ARC-e, ARC-c[8], HellaSwag[36], which are commonly used for evaluating the performance of LLM quantization methods. To further enhance evaluation coverage, we also include COPA[27] for causal reasoning and LAMBADA[25] for long-context language modeling. All evaluations are conducted using the open-source LLM evaluation framework, LM-Evaluation-Harness[24].

**Details of Experiments.** All experiments are conducted with PyTorch on NVIDIA GeForce RTX 3090 GPUs with 24GB of memory. For the calibration data, we follow the settings adopted in GPTQ and BiLLM, selecting 128 samples from the C4 dataset, with a sequence length of 2048. During quantization, we set the block size to 128 in BiLLM, PB-LLM, ARB-LLM, and HBLLM.

Baselines. We compare HBLLM against several state-of-the-art LLM binarization methods, including BiLLM, ARB-LLM and PB-LLM, ensuring that all implementations adhere to the details provided in their respective papers. BiLLM, ARB-LLM and PB-LLM all utilize the PTQ approach for model calibration through OBQ based method of GPTQ. For ARB-LLM, we evaluate two of its best-performing variants, ARB-LLM<sub>X</sub> and ARB-LLM<sub>RC</sub>. Both ARB-LLM<sub>X</sub> and ARB-LLM<sub>RC</sub> employ the salient column bitmap and group bitmap (CGB) for better performance. For PB-LLM, which allows variable ratios of salient weights to enhance accuracy, we have set the ratio of salient weights to 10% to ensure the average bit width of weight parameters remains below 2 bits. Given the significant accuracy improvements demonstrated by HBLLM over traditional binarization techniques, we also include a comparison with a leading method using orthogonal transforms: FrameQuant. For FrameQuant, quantization is performed not in the original weight space but in the structured orthogonal basis constructed through Fusion Frames. We evaluate two configurations: FrameQuant (r=1.0) and FrameQuant (r=1.1), where the redundancy factor r controls the amount of redundancy introduced during the transformation.

# 4.2 Perplexity and Accuracy Results of 1-2 Bit Quantized Models

The perplexity and zero-shot accuracy results of previous 1-2 bit quantization methods and the proposed HBLLM are presented in Table1. HBLLM consistently outperforms existing 1-2 bit quantization techniques across all evaluation metrics.

Specifically, HBLLM reduces the language modeling perplexity by 33%-66% compared to previous methods, while achieving substantial improvements in QA task accuracy, with relative gains ranging from -0.73% to +11.3%. In our experiments, HBLLM slightly outperforms FrameQuant, a 2.2-bit quantization method, and exhibits a particularly significant advantage on the LLaMA-3-8B model. Moreover, when compared with BiLLM and ARB-LLM<sub>X</sub>, HBLLM-col, demonstrates a clear advantage in both perplexity and accuracy, despite operating at comparable or lower bit-widths. These results indicate that HBLLM effectively narrows the performance gap between quantized models and their Float16 counterparts, achieving  $1.22\times$  to  $2.48\times$  of the original perplexity and retaining 73.8%-88.8% of the original QA accuracy.

# 4.3 Ablation Study

Salient Column Selection Criterion. To evaluate the impact of selection criteria in salient column screening on quantization effectiveness, we compare two strategies: the column  $\ell_1$  norm and the column  $\ell_2$  norm as significance indicators. Experimental results in Table 2a reveal that the column  $\ell_2$  norm consistently achieves lower quantization error and superior performance in downstream tasks, indicating its greater effectiveness in capturing energy distribution across columns and enhancing quantization quality.

Table 2: Ablation study on LLaMA2-7B. Results are measured by perplexity, with final results highlighted in **bold**.

(a) Study of salient column selection criterion

(1 \ C) 1	c		1
(b) Study	ΩŤ	graiining	oraniilarity
(b) Study	O1	Stouping	granularity

Method	Selection criterion	Wiki2↓	PTB↓
HBLLM-row	$\ell_1 \\ \ell_2$	10.78 <b>10.52</b>	143.7 <b>89.23</b>
HBLLM-col	$\ell_1$ $\ell_2$	11.45 <b>11.33</b>	308.2 <b>150.6</b>

Method	Group Partition	Wiki2↓	PTB↓
HBLLM-row	global	16.32	1990
	row-wise	<b>11.08</b>	<b>95.58</b>
HBLLM-col	global	13.99	1546
	row-wise	<b>12.02</b>	<b>146.1</b>

(c) Effectiveness of shared mean

(d) Study of partitioning candidates number

Method	Shared mean Wiki2↓		PTB↓	
HBLLM-row	×	11.08 <b>10.52</b>	95.58 <b>89.23</b>	
HBLLM-col	X ./	12.02 <b>11.33</b>	<b>146.1</b> 150.6	

Method	Candidate number	Wiki2↓	РТВ↓
HBLLM-row	10 20	11.16 11.32	108.8 165.8
TIBEENT TOW	40 80	<b>11.08</b> 11.13	<b>95.58</b> 113.8

**Granularity of Group Quantization.** To explore the influence of grouping granularity on model performance, we compare global grouping with row-wise grouping strategies, evaluating both quantization error and perplexity, as shown in Table 2b. The results reveal that row-wise grouping significantly reduces quantization error and achieves lower perplexity compared to global grouping. This suggests that finer-grained row-wise partitioning better preserves local data fidelity, leading to improved quantized inference performance.

**Shared Mean Strategy.** Under the standard dual-partition quantization setting, we further explore a compression strategy that shares the quantization center across two partitions within each row. By unifying the mean for both partitions, the storage overhead of quantization coefficients can be significantly reduced. Experimental results in Table 2c demonstrate that the shared mean strategy even slightly reduces quantization error without degrading perplexity, verifying its effectiveness and practicality.

Choice of Partitioning Number. We investigate the impact of varying the number of partition candidates on final quantization performance under the row-wise grouping setting. Specifically, for each row, we generate partition candidates based on absolute value percentiles ranging from 10% to 90%, and evaluate the corresponding quantization error and perplexity, as shown in Table 2d. Experimental results indicate that moderately increasing the number of partition candidates can effectively reduce quantization error and further lower perplexity, while excessive partitioning yields diminishing returns and increases computational cost. Consequently, we adopt 40 partition candidates as the default setting to balance performance and efficiency.

# 4.4 Time and Memory Analysis

Time Comparison. As a binary PTQ framework, HBLLM eliminates the need for finetuning. The introduction of Haar wavelet transforms requires additional computation during quantization, yet this overhead remains fully acceptable. As shown in Table 3, HBLLM increases the quantization time by approximately 20%-30% compared to BiLLM across different model sizes. It is worth noting that ARB-LLM<sub>X</sub> and FrameQuant fail to complete quantization for LLaMA-1-13B and LLaMA-1-30B under the single-GPU-24 GB setting, while HBLLM

Table 3: Time comparison between LLM binarization methods and our HBLLM on LLaMA-1 with different model sizes.

Method	<b>7B</b>	13B	30B
BiLLM	36min	71min	142min
$ARB-LLM_x$	88min	Х	Х
$ARB-LLM_{RC}$	76min	119min	239min
PB-LLM	18min	29min	57min
FrameQuant	14min	22min	X
HBLLM	44min	98min	173min

successfully completes the process, demonstrating better scalability.

Memory Comparison. As shown in Table 4, HBLLM-col achieves better performance while occupying a storage size comparable to ARB-LLM. By employing a grouped shared-mean strategy, HBLLM improves compression efficiency without sacrificing performance. Specifically, HBLLM-col applies Haar transforms along the column dimension, such that only one grouped quantization operation is required per row on the transformed coefficients. Compared to HBLLM-row, this leads to reduced data fidelity but provides clear advantages in storage cost. The detailed storage calculation formulas can be found in the appendix D.

**Memory Comparison.** As shown in Table 4, Table 4: Memory comparison LLM binarization HBLLM-col achieves better performance while occupying a storage size comparable to ARB-ferent model sizes.

Method	7B	13B
FP16	13.48GB	26.03GB
BiLLM	2.93GB	5.36GB
$ARB-LLM_x$	3.23GB	5.95GB
$ARB-LLM_{RC}$	2.83GB	5.17GB
PB-LLM	2.91GB	5.33GB
FrameQuant	11.36GB	16.08GB
HBLLM-row	3.09GB	5.89GB
HBLLM-col	2.67GB	5.06GB

# 4.5 Inference Latency Estimation

To evaluate the inference latency of HBLLM, we conduct an experiment that combines direct measurement with estimation. Due to there is no existing inference framework that fully supports the dequantization algorithm used in HBLLM, we test GEMV on layers from the OPT-175B model instead. The tests are run on an NVIDIA P100 GPU following the GPTQ benchmark setup [1]. Our estimation results show that the inference latency of HBLLM is approximately 31.8% of the FP16 baseline inference time. For more details, please refer to the appendix G.

#### 5 Conclusion

We introduce a 1-bit weight only quantization HBLLM, which applies Haar transform to BILLM pipeline. Besides quantifying the coefficients on frequence domain, HBLLM integrates two innovative structure-aware grouping strategies to enhance fidelity. Furthermore, HBLLM optimize storage efficiency. As a results, HBLLM outperforms SOTA QAT quantization methods of LLM at 1-bit across different LLM families and tests. The current HBLLM supports only quantized dense models. Next, we will focus on the MoE PTQ algorithm.

# Acknowledgements

We thank all constructive comments from anonymous reviewers. This work is partially supported by the National Key Research and Development Program of China under Grant No.2023YFB3001704.

# References

- [1] Adepu H, Zeng Z, Zhang L, et al. FrameQuant: flexible low-bit quantization for transformers. *Proceedings of the 41st International Conference on Machine Learning*. 2024: 203–227.
- [2] Ashkboos S, Croci M L, do Nascimento M G, et al. SliceGPT: Compress Large Language Models by Deleting Rows and Columns. *The Twelfth International Conference on Learning Representations*. 2024.
- [3] Lin J, Tang J, Tang H, et al. AWQ: Activation-aware weight quantization for on-device LLM compression and acceleration. *Proceedings of Machine Learning and Systems*, 2024, 6: 87–100.
- [4] Bisk Y, Zellers R, Gao J, et al. PIQA: Reasoning about physical commonsense in natural language. *Proceedings of the AAAI Conference on Artificial Intelligence*. 2020: 7432–7439.
- [5] Chee J, Cai Y, Kuleshov V, et al. QuIP: 2-bit quantization of large language models with guarantees. *Advances in Neural Information Processing Systems*, 2023, 36: 4396–4429.
- [6] Chen H, Lv C, Ding L, et al. DB-LLM: Accurate dual-binarization for efficient LLMs. Findings of the Association for Computational Linguistics: ACL 2024. 2024: 8719–8730.
- [7] Clark C, Lee K, Chang M W, et al. BoolQ: Exploring the surprising difficulty of natural yes/no questions. *Proceedings of NAACL-HLT*. 2019: 2924–2936.

- [8] Clark P, Cowhey I, Etzioni O, et al. Think you have solved question answering? Try ARC, the AI2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- [9] Duan Y, Liu F, Jiao L, et al. SAR image segmentation based on convolutional-wavelet neural network and Markov random field. *Pattern Recognition*, 2017, 64: 255–267.
- [10] Edalati A, Ghaffari A, Nejad M G, et al. OAC: Output-adaptive calibration for accurate post-training quantization. Proceedings of the AAAI Conference on Artificial Intelligence. 2025, 39(16): 16453–16461.
- [11] Frantar E, Alistarh D. SparseGPT: Massive language models can be accurately pruned in one-shot. *International Conference on Machine Learning*. 2023: 10323–10337.
- [12] Frantar E, Ashkboos S, Hoefler T, et al. OPTQ: Accurate post-training quantization for generative pre-trained transformers. 11th International Conference on Learning Representations. 2023.
- [13] Gong R, Ding Y, Wang Z, et al. A survey of low-bit large language models: Basics, systems, and algorithms. *arXiv* preprint arXiv:2409.16694, 2024.
- [14] Huang H, He R, Sun Z, et al. Wavelet-SRNet: A wavelet-based CNN for multi-scale face super resolution. Proceedings of the IEEE International Conference on Computer Vision. 2017: 1689–1697.
- [15] Huang W, Liu Y, Qin H, et al. BiLLM: Pushing the limit of post-training quantization for LLMs. *Proceedings of the 41st International Conference on Machine Learning*. 2024: 20023–20042.
- [16] Huang W, Zheng X, Ma X, et al. An empirical study of LLaMA3 quantization: From LLMs to MLLMs. *Visual Intelligence*, 2024, 2(1): 36.
- [17] Jo D, Kim T, Kim Y. Mixture of scales: Memory-efficient token-adaptive binarization for large language models. Advances in Neural Information Processing Systems, 2024, 37: 137474–137494.
- [18] Li Z, Yan X, Zhang T, et al. Arb-LLM: Alternating refined binarizations for large language models. *arXiv* preprint arXiv:2410.03129, 2024.
- [19] Liang C, Zuo S, Zhang Q, et al. Less is more: Task-aware layer-wise distillation for language model compression. *International Conference on Machine Learning*. 2023: 20852–20867.
- [20] Mallat S. A Wavelet Tour of Signal Processing. Elsevier, 1999.
- [21] Marcus M, Kim G, Marcinkiewicz M A, et al. The Penn Treebank: Annotating predicate argument structure. *Human Language Technology: Proceedings of a Workshop held at Plainsboro*. 1994.
- [22] Merity S, Xiong C, Bradbury J, et al. Pointer Sentinel Mixture Models. *International Conference on Learning Representations*. 2017.
- [23] Mihaylov T, Clark P, Khot T, et al. Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. 2018: 2381–2391.
- [24] Mukherjee S, Liu X, Zheng G, et al. Few-shot learning evaluation in natural language understanding. Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track. 2021.
- [25] Paperno D, Kruszewski G, Lazaridou A, et al. The LAMBADA dataset: Word prediction requiring a broad discourse context. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. 2016: 1525–1534.
- [26] Raffel C, Shazeer N, Roberts A, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 2020, 21(140): 1–67.
- [27] Roemmele M, Bejan C A, Gordon A S. Choice of Plausible Alternatives: An Evaluation of Commonsense Causal Reasoning. *AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*. 2011: 90–95.
- [28] Sakaguchi K, Bras R L, Bhagavatula C, et al. Winogrande: An adversarial Winograd schema challenge at scale. *Communications of the ACM*, 2021, 64(9): 99–106.
- [29] Shang Y, Yuan Z, Wu Q, et al. PB-LLM: Partially binarized large language models. *The Twelfth International Conference on Learning Representations*. 2024.
- [30] Shridhar K, Stolfo A, Sachan M. Distilling reasoning capabilities into smaller language models. Findings of the Association for Computational Linguistics: ACL 2023, 2023: 7059–7073.

- [31] Sun M, Liu Z, Bair A, et al. A simple and effective pruning approach for large language models. *arXiv* preprint arXiv:2306.11695, 2023.
- [32] Touvron H, Lavril T, Izacard G, et al. LLaMA: Open and efficient foundation language models. *arXiv* preprint arXiv:2302.13971, 2023.
- [33] Xiao G, Lin J, Seznec M, et al. SmoothQuant: Accurate and efficient post-training quantization for large language models. *International Conference on Machine Learning*. 2023: 38087–38099.
- [34] Xu Y, Han X, Yang Z, et al. OneBit: Towards extremely low-bit large language models. *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2024: 66357–66382.
- [35] Yao Z, Yazdani Aminabadi R, Zhang M, et al. ZeroQuant: Efficient and affordable post-training quantization for large-scale transformers. Advances in Neural Information Processing Systems, 2022, 35: 27168–27183.
- [36] Zellers R, Holtzman A, Bisk Y, et al. HellaSwag: Can a machine really finish your sentence? *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2019: 4791–4800.
- [37] Zhang S, Roller S, Goyal N, et al. OPT: Open pre-trained transformer language models. *arXiv preprint* arXiv:2205.01068, 2022.

# **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: The abstract and introduction clearly state our main contributions, which are supported by the theoretical and empirical results presented in the Section 3 and Section 4.

#### 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: Please refer to Section 5.

- 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: [Yes]

Justification: Please refer to Appendix.

#### 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: Please refer to Section 4.

- 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: [Yes]

Justification: The code will be made publicly available upon acceptance of this paper.

#### 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, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Please refer to Section 4.1.

#### 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: [No]

Justification: Error bars are not reported because it would be too computationally expensive.

#### 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: Quantization for models <30B was run on  $4\times$ RTX 3090 (24GB), and for models  $\geq$ 30B on A800-80GB. See Section 4.1.

#### 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 carefully read and follow the NeurIPS 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: [No]

Justification: The paper is purely fundamental research and does not involve social impact.

- 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: [No]

Justification: This paper does not release new models or datasets.

#### 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: Please refer to Appendix.

#### 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 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: The paper does not release new assets.

#### 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: : The paper does not involve crowdsourcing nor research with human subjects.

#### 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: The paper does not involve crowdsourcing nor research with human subjects.

#### 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.

#### 16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: LLMs were not involved as core components in the development of the method or results in this paper.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.