



Quant-dLLM: Post-Training Extreme Low-Bit Quantization for Diffusion Large Language Models

Tianao Zhang^{1,2*}, Zhiteng Li^{1*}, Xianglong Yan¹, Haotong Qin³, Yong Guo⁴, Yulun Zhang^{1†}

¹Shanghai Jiao Tong University, ²Zhiyuan College, SJTU, ³ETH Zürich, ⁴Huawei



Symmetric and Asymmetric Binarization

$$\text{sign}(x) = \begin{cases} +1, & \text{if } x \geq 0, \\ -1, & \text{others.} \end{cases} \begin{array}{l} \xrightarrow{\text{sym.}} \mathbf{W}_b = \alpha \cdot \text{sign}(\mathbf{W}_f) \\ \xrightarrow{\text{asym.}} \mathbf{W}_b = \alpha \cdot \text{sign}(\mathbf{W}_f - \mu) + \mu \end{array}$$

- Typically, we define $\mu = \text{mean}(\mathbf{W}_f)$.
- Let $\mathbf{B} = \text{sign}(\mathbf{W}_f - \mu)$, $\tilde{\mathbf{W}} = \mathbf{W}_f - \mu$.
- **Objective of Binarization:**

$$\arg \min_{\alpha, \mathbf{B}} \|\tilde{\mathbf{W}} - \alpha \mathbf{B}\|_F^2$$

- Least square solution for the scaling factor:

$$\alpha = (\mathbf{B}^\top \mathbf{B})^{-1} \mathbf{B}^\top \tilde{\mathbf{W}}$$

- Typically, for [row-wise binarization](#) (also referred to as [per-channel binarization](#)), we have

$$\alpha = \frac{1}{m} \mathbf{B}^\top \tilde{\mathbf{W}}, \quad \text{where } \mathbf{B} \in \mathbb{R}^m$$



Binarization alone row-column axes

$$\text{sign}(x) = \begin{cases} +1, & \text{if } x \geq 0, \\ -1, & \text{others.} \end{cases} \quad \mathbf{W}_b = \alpha_r \alpha_c^T \cdot \text{sign}(\mathbf{W}_f)$$

- Here, we assume symmetric quantization, so we define $\mu = 0$.
- Let $\mathbf{B} = \text{sign}(\mathbf{W}_f)$.
- **Objective of Binarization:**

$$\arg \min_{\alpha_r, \alpha_c, \mathbf{B}} \|\tilde{\mathbf{W}} - (\alpha_r \alpha_c^T) \odot \mathbf{B}\|_F^2$$

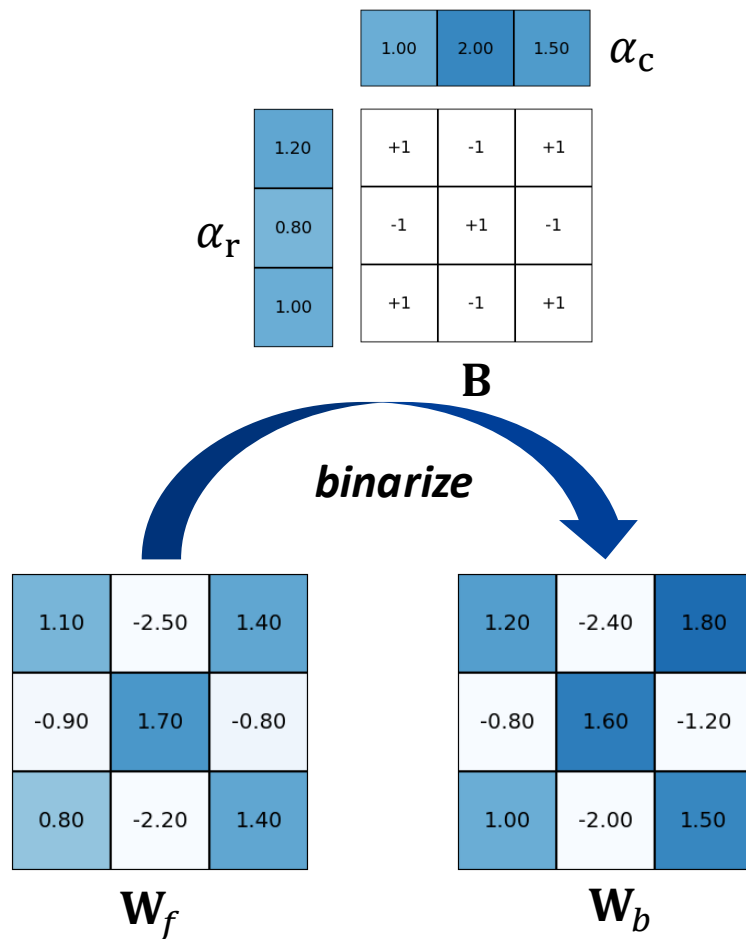
- Typically, for Binarization alone row-column axes, we have

$$\alpha_r = \frac{1}{m} \mathbf{B}^T \tilde{\mathbf{W}}, \quad \text{where } \mathbf{B} \in \mathbb{R}^m$$
$$\alpha_c = \frac{1}{n} \mathbf{B}^T (\text{diag}(\alpha_r)^{-1} \tilde{\mathbf{W}}), \quad \text{where } \mathbf{B} \in \mathbb{R}^n$$

Binarization / INT1 Quantization



A toy example



$R \neq 0$, Higher-order Binarization is required to further minimize the quantization error!





1st-order and knd-order Binarization

- **Definition**

- **1st-order** binarization: $\hat{\mathbf{W}} = (\alpha_r \alpha_c^T) \odot \mathbf{B}$
- **knd-order** binarization: $\hat{\mathbf{W}} = \sum (\alpha_{r,k} \alpha_{c,k}^T) \odot \mathbf{B}_k$

*For outliers, it's more accurate
but uses more memory.*

- **Optimization**

- **1st-order** binarization: $\mathbf{B} = \text{sign}(\mathbf{W}_f)$, $\alpha_r = \frac{1}{m} \mathbf{B}^T \tilde{\mathbf{W}}$, and $\alpha_c = \frac{1}{n} \mathbf{B}^T (\text{diag}(\alpha_r)^{-1} \tilde{\mathbf{W}})$
- **knd-order** binarization: how to obtain the optimal $\alpha_{r,1}, \alpha_{r,k}, \alpha_{c,1}, \alpha_{c,k}, \mathbf{B}_1$, and \mathbf{B}_k ?

- **Objective of Binary Residual Approximation:**

$$\arg \min_{\alpha_{r,1}, \alpha_{c,1}, \mathbf{B}_1} \|\mathbf{W} - (\alpha_{r,1} \alpha_{c,1}^T) \odot \mathbf{B}_1\|_F^2, \quad \arg \min_{\alpha_{r,k}, \alpha_{c,k}, \mathbf{B}_k} \|\mathbf{R}_{k-1} - (\alpha_{r,k} \alpha_{c,k}^T) \odot \mathbf{B}_k\|_F^2$$

$$\text{where } \mathbf{R}_0 = \mathbf{W}, \quad \mathbf{R}_k = \mathbf{R}_{k-1} - (\alpha_{r,k} \alpha_{c,k}^T) \odot \mathbf{B}_k.$$

LLaDA

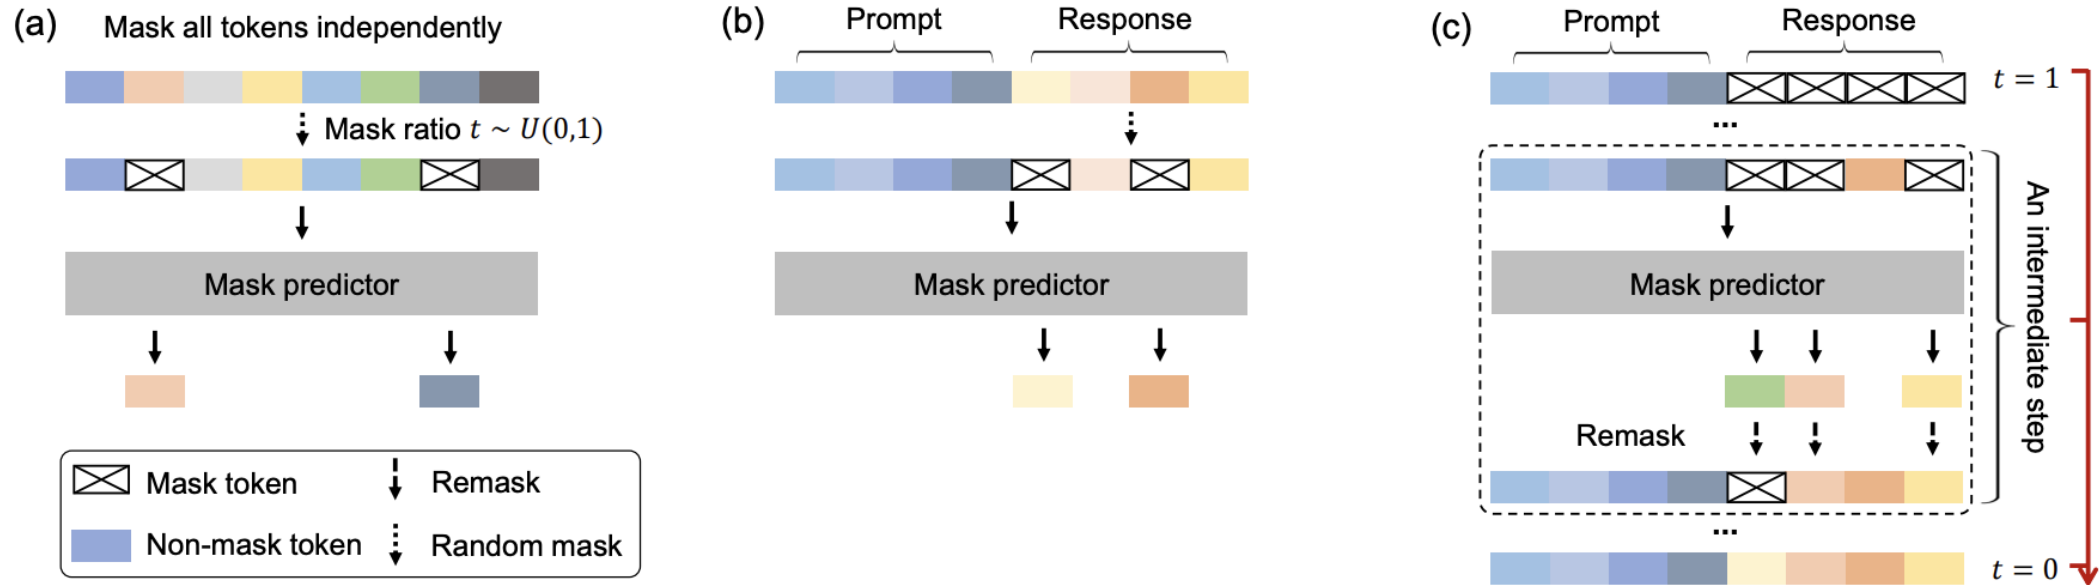
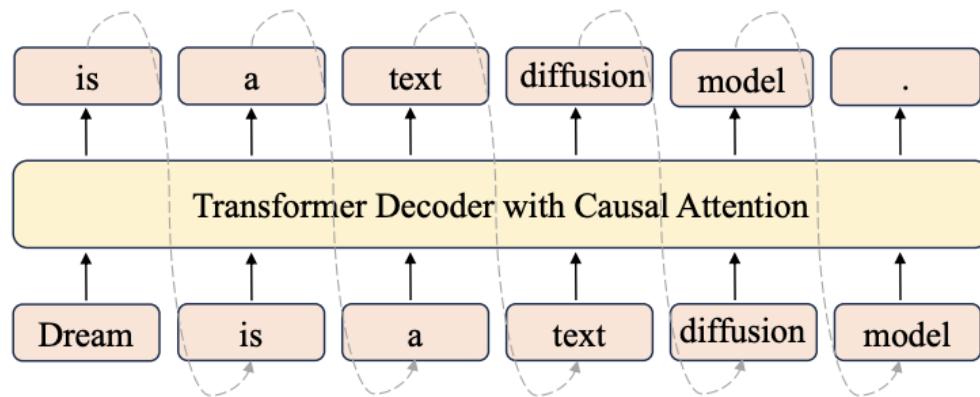
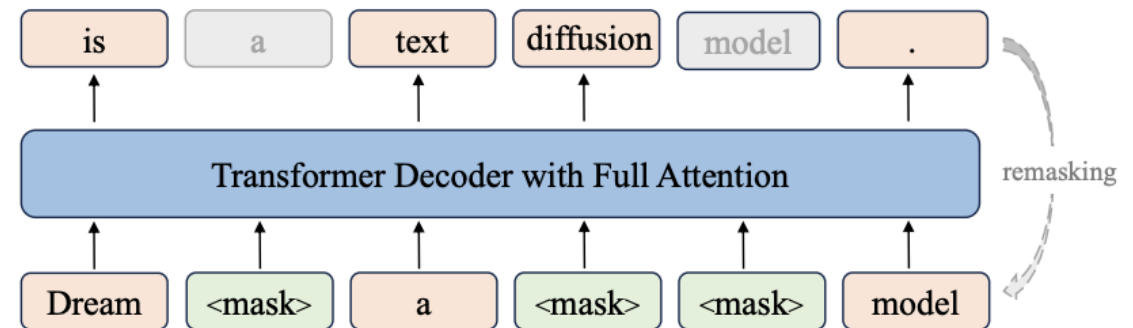


Figure 2: **Overview of LLaDA.** (a) Pre-training. LLaDA is trained on text with random masks applied independently to all tokens at the same ratio $t \sim U[0, 1]$. (b) SFT. Only response tokens are possibly masked. (c) Sampling. LLaDA simulates a diffusion process from $t = 1$ (fully masked) to $t = 0$ (unmasked), predicting all masks simultaneously at each step with flexible remask strategies.

Dream



(a) Autoregressive Modeling



(b) Diffusion Modeling in Dream

Figure 2: Comparison of autoregressive modeling and diffusion modeling in Dream. Dream predicts all the masked tokens in a shifted manner, allowing for maximum architectural alignment and weight initialization with AR models.

Observation

1. Static calibration creates a **distribution mismatch** with the dynamic masked activations of diffusion LLMs.
2. The optimization of binarization parameters α_r and α_c does not adequately reflect the real-world scenario, as it does not account for the **input data X**.
3. Uniform bit-width allocation ignores the **column-wise clustering** of highly salient weights.

Solution

1. **Masked Calibration Simulation (MCS)**: To reduce the distribution mismatch, MCS generates timestep-aware, partially visible calibration inputs that mirror the model's native diffusion denoising process.
2. **Data-aware Any-order Quantizer (DAQ)**: To incorporate input data **X**, DAQ optimizes a data-aware objective that directly minimizes output activation error.
3. **Adaptive Blockwise Mixed Precision (ABMP)**: To address uneven salience, ABMP calculates block-wise importance scores to dynamically assign mixed precision.

Overview

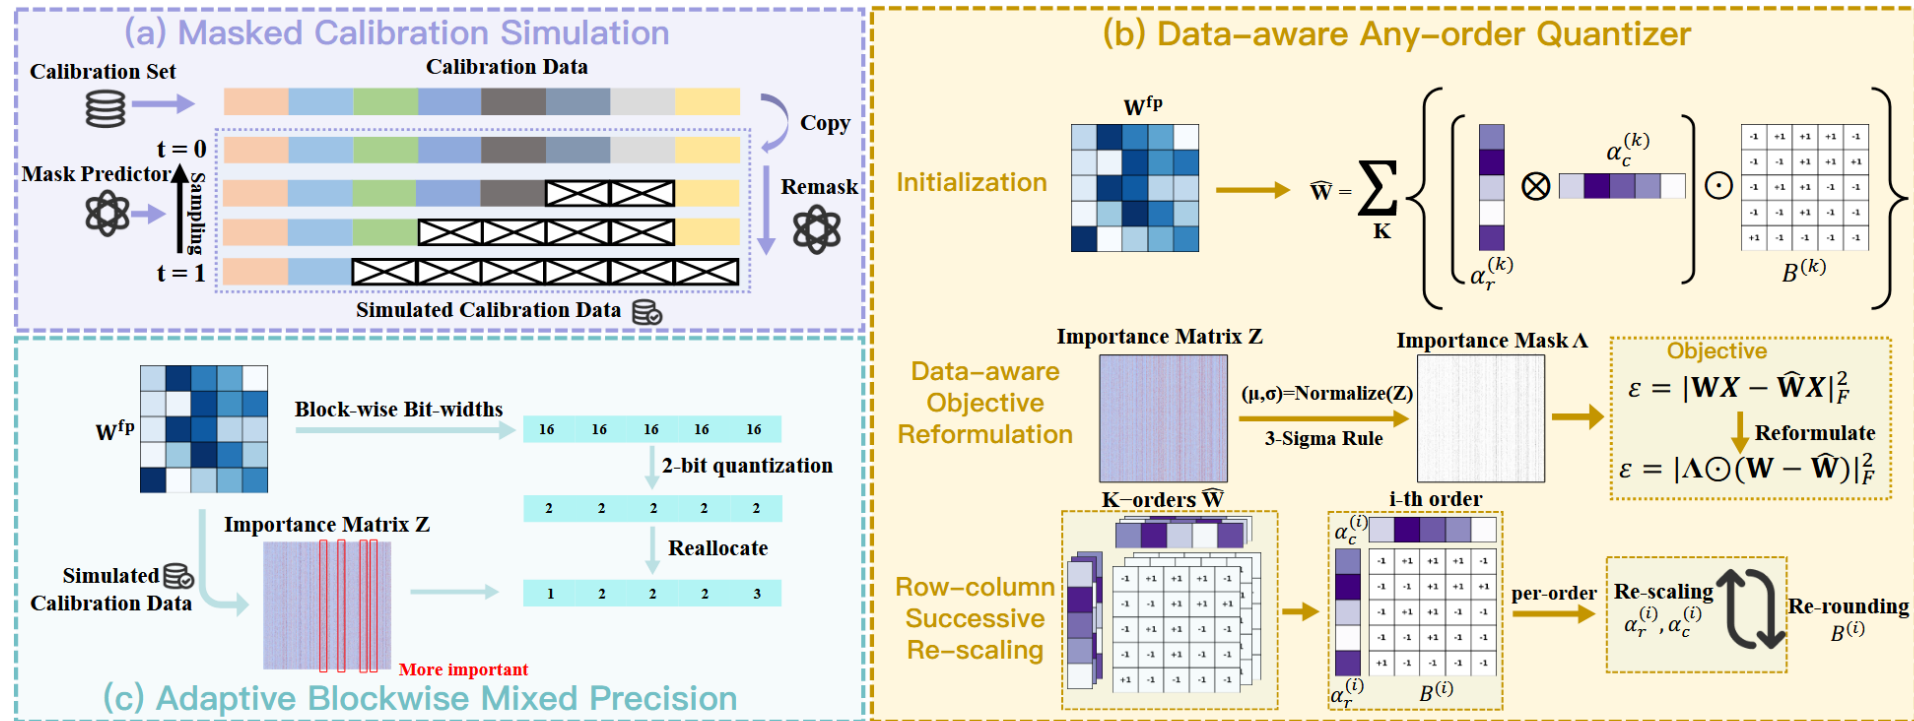


Figure 2: Overview of our Quant-dLLM. **Masked Calibration Simulation**: Aligns calibration with diffusion by simulating masked, timestep-aware inputs. **Adaptive Blockwise Mixed Precision**: Assigns binary orders by importance under a 2-bit average. **Data-aware Any-order Quantizer**: Builds multi-binary RC forms with data-aware optimization.

Masked Calibration Simulation (MCS)



Conventional PTQ calibration assumes:

- Autoregressive activations.
- Fully visible tokens.
- Static token distribution.

distribution mismatch

In diffusion LLM:

- Token visibility depends on timestep.
- Activation distribution changes over denoising steps.



simulate the timestep-dependent masking pattern

Masked Calibration Simulation (MCS)

- Pseudocode of Masked Calibration Simulation

Algorithm 1 MCS: Masked Calibration Simulation

```
func MCS( $\mathcal{X}, T, \gamma, \alpha(t)$ )
Input:  $\mathcal{X}$  — full sequences;  $T$  — total timesteps;
 $\gamma$  — prefix ratio;  $\alpha(t)$  — visibility schedule
Output: Masked calibration set  $\mathcal{D}_{\text{calib}}$ 
1:  $\mathcal{D}_{\text{calib}} \leftarrow \emptyset$ 
2: for all  $x \in \mathcal{X}$  do
3:    $L \leftarrow$  length of  $x$ 
4:    $\mathcal{P} \leftarrow \{1, \dots, \lfloor \gamma \cdot L \rfloor\}$ 
5:   for  $t = 1$  to  $T$  do
6:      $\alpha \leftarrow \alpha(t)$ 
7:     for  $i = 1$  to  $L$  do
8:        $r_i \leftarrow 1$  if  $i \in \mathcal{P}$  else Bernoulli( $\alpha$ )
9:        $\tilde{x}_i(t) \leftarrow x_i$  if  $r_i = 1$  else [MASK]
10:    end for
11:     $\mathcal{D}_{\text{calib}} \leftarrow \mathcal{D}_{\text{calib}} \cup \{\tilde{x}(t)\}$ 
12:  end for
13: end for
14: return  $\mathcal{D}_{\text{calib}}$ 
```

Data-aware Any-order Quantizer (DAQ)

- **K^{nd} -order binarization**

- k^{nd} -order binarization: $\widehat{\mathbf{W}} = \sum (\alpha_{r,k} \alpha_{c,k}^T) \odot \mathbf{B}_k$

- **Quantization Error**

- **weight-activation:** $\mathcal{L}_2 = \|\mathbf{W}\mathbf{X} - \widehat{\mathbf{W}}\mathbf{X}\|_F^2$, where $\mathbf{X} \in \mathbb{R}^{B \times L \times m}$ is a small [calibration set](#).


- **Reformulation of \mathcal{L}_2**

- Let $\mathbf{S} = \sum_b \mathbf{X}_b^T \mathbf{X}_b$ and $\mathbf{R} = \mathbf{W} - \widehat{\mathbf{W}}$, where $\mathbf{S} \in \mathbb{R}^{m \times m}$.

- Then $\mathcal{L}_2 = \|\mathbf{W}\mathbf{X} - \widehat{\mathbf{W}}\mathbf{X}\|_F^2 = \langle \mathbf{S}, \mathbf{R}^T \mathbf{R} \rangle_F = \text{Tr}(\mathbf{R}\mathbf{S}\mathbf{R}^T)$.

- **Consider MCS**

- Then $\mathcal{L}_2 = \text{Tr}((\mathbf{W} - \widehat{\mathbf{W}})\mathbf{S}_{\text{MCS}}(\mathbf{W} - \widehat{\mathbf{W}})^T)$.



Data-aware Any-order Quantizer (DAQ)

- **Data-aware Objective Reformulation (DOR)**

- Quantization error is not uniformly distributed but concentrates on a small, critical subset of weights.
- Importance matrix $\mathbf{Z} \in \mathbb{R}^{n \times m} = (\mathbf{W}\mathbf{D}^{-1}) \odot (\mathbf{W}\mathbf{D}^{-1})$, where $\mathbf{D} = \text{diag}((\mathbf{S}_{\text{MCS}} + \gamma\mathbf{I})^{-1})$.
- Then standardize $\tilde{\mathbf{Z}} = (\mathbf{Z} - \mu) \oslash \sigma$.
- Importance mask $\mathbf{\Lambda} = \mathbf{1} + (\lambda - 1)\mathbf{\Pi}$, where $\mathbf{\Pi} = \mathbf{I}(|\tilde{\mathbf{Z}}| > 3)$.

- **Reformulation of \mathcal{L}_{Λ}**

- $$\mathcal{L}_{\Lambda} = \|\mathbf{\Lambda} \odot (\mathbf{W} - \widehat{\mathbf{W}})\|_F^2 = \|\mathbf{\Lambda} \odot (\mathbf{W} - \sum(\alpha_{r,k}\alpha_{c,k}^T) \odot \mathbf{B}_k)\|_F^2.$$



Data-aware Any-order Quantizer (DAQ)


- Row-column Successive Re-scaling (RSR)

- ① Initialize

$$\mathbf{B} := \text{sign}(\mathbf{W})$$


$$\alpha_r = \frac{1}{m} \mathbf{B}^\top \tilde{\mathbf{W}}, \text{ where } \mathbf{B} \in \mathbb{R}^m$$

$$\alpha_c = \frac{1}{n} \mathbf{B}^\top (\text{diag}(\alpha_r)^{-1} \tilde{\mathbf{W}}), \text{ where } \mathbf{B} \in \mathbb{R}^n \quad \mathcal{L}_\Lambda = \|\Lambda \odot (\mathbf{W} - (\alpha_r \alpha_c^\top) \odot \mathbf{B})\|_F^2$$

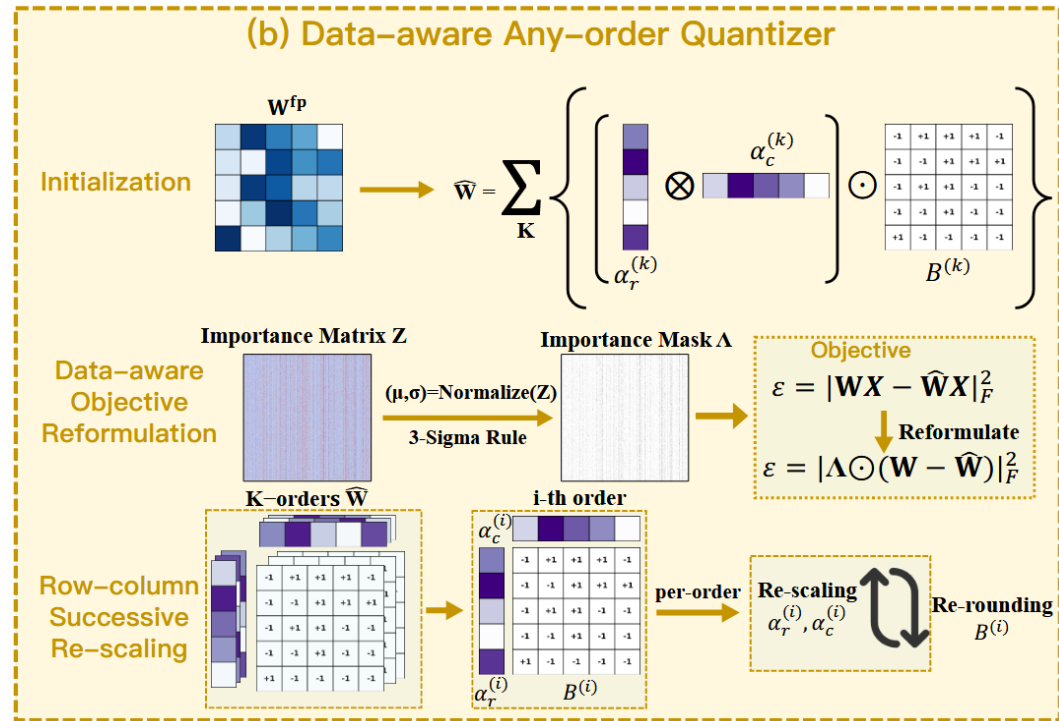


$$\alpha_r \leftarrow \frac{(\Lambda^2 \odot \mathbf{W} \odot \mathbf{B}) \alpha_c}{\Lambda^2 (\text{diag}(\alpha_c) \alpha_c) + \varepsilon \mathbf{I}}$$

$$\alpha_c \leftarrow \frac{(\Lambda^2 \odot \mathbf{W} \odot \mathbf{B})^\top \alpha_r}{\Lambda^2 (\text{diag}(\alpha_r) \alpha_r) + \varepsilon \mathbf{I}}$$



Data-aware Any-order Quantizer (DAQ)



Algorithm 2 DAQ: Data-aware any-order Quantizer

```

func DAQ(W, Z, K, T, λ)
Input W, Z ∈ ℝn×m, K, T ∈ ℕ, λ > 1
Output
     $\hat{W} = \sum_{k=1}^K \left( \alpha_r^{(k)} \alpha_c^{(k)\top} \right) \odot B^{(k)}$ 
    1:  $\Lambda := \text{build\_importance\_mask}(Z, \lambda)$ 
    2:  $\hat{W} := \mathbf{0}_{n \times m}$ 
    3: for  $k = 1, 2, \dots, K$  do
    4:    $\mathbf{R} \leftarrow W - \hat{W}$ 
    5:    $\alpha_r^{(k)}, \alpha_c^{(k)}, B^{(k)} \leftarrow \text{binary\_rc\_init}(\mathbf{R})$ 
    6:    $\mathbf{S}^{(k)} \leftarrow \alpha_r^{(k)} (\alpha_c^{(k)})^\top$ 
    7:    $\hat{W} \leftarrow \hat{W} + \mathbf{S}^{(k)} \odot B^{(k)}$ 
    8: end for
    9: for  $t = 1, 2, \dots, T$  do
    10:  for  $k = 1, 2, \dots, K$  do
    11:     $\hat{W}_{\setminus k} \leftarrow \sum_{q \neq k} (\mathbf{S}^{(q)} \odot B^{(q)})$ 
    12:     $\mathbf{R}^{(k)} \leftarrow W - \hat{W}_{\setminus k}$ 
    13:     $\alpha_r^{(k)} \leftarrow \text{update\_}\alpha_r(\mathbf{R}^{(k)}, B^{(k)}, \alpha_c^{(k)}, \Lambda)$ 
    14:     $\alpha_c^{(k)} \leftarrow \text{update\_}\alpha_c(\mathbf{R}^{(k)}, B^{(k)}, \alpha_r^{(k)}, \Lambda)$ 
    15:     $\mathbf{S}^{(k)} \leftarrow \alpha_r^{(k)} (\alpha_c^{(k)})^\top$ 
    16:  end for
    17:   $\{B^{(k)}\}_{k=1}^K \leftarrow \text{update\_B}(W, \{\alpha_r^{(k)}\}_{k=1}^K, \{\alpha_c^{(k)}\}_{k=1}^K)$ 
    18:   $\hat{W} \leftarrow \sum_{k=1}^K \mathbf{S}^{(k)} \odot B^{(k)}$ 
    19: end for
    20: return  $\hat{W}$ 

func binary_rc_init(X)
  1:  $\alpha_r \leftarrow \frac{1}{m} |X| \mathbf{1}_m$ 
  2:  $\alpha_c \leftarrow \frac{1}{n} |X|^\top \text{diag}(\alpha_r)^{-1} \mathbf{1}_n$ 
  3:  $B \leftarrow \text{sign}(X)$ 
  4: return  $\alpha_r, \alpha_c, B$ 

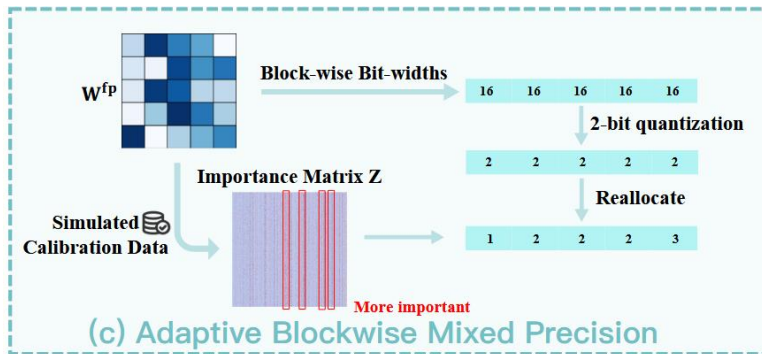
func update_αr(X, B, αc, Λ)
  1:  $u \leftarrow (\Lambda^2 \odot X \odot B) \alpha_c$ 
  2:  $v \leftarrow (\Lambda^2 (\text{diag} \alpha_c) \alpha_c) + \varepsilon \mathbf{1}_n$ 
  3: return  $u \oslash v$ 

func update_αc(X, B, αr, Λ)
  1:  $u \leftarrow (\Lambda^2 \odot X \odot B)^\top \alpha_r$ 
  2:  $v \leftarrow (\Lambda^2)^\top (\text{diag}(\alpha_r) \alpha_r) + \varepsilon \mathbf{1}_m$ 
  3: return  $u \oslash v$ 

func update_B(W, {αr(k)}k=1K, {αc(k)}k=1K)
  1: for  $i = 1, \dots, n$  do
  2:   for  $j = 1, \dots, m$  do
  3:      $s \leftarrow [(\alpha_r^{(k)})_i (\alpha_c^{(k)})_j]_{k=1}^K$ 
  4:      $b \leftarrow \text{search}(W_{ij}, s)$ 
  5:     for  $k = 1, \dots, K$  do
  6:        $(B^{(k)})_{ij} \leftarrow b_k$ 
  7:     end for
  8:   end for
  9: end for
  10: return  $\{B^{(k)}\}_{k=1}^K$ 

func build_importance_mask(Z, λ)
  1:  $\mu \leftarrow \text{mean}(Z), \sigma \leftarrow \text{std}(Z)$ 
  2:  $M \leftarrow \mathbb{I}(|Z - \mu| > 3\sigma)$ 
  3: return  $\mathbf{1}_{n \times m} + (\lambda - 1)M$ 
  
```

Adaptive Blockwise Mixed Precision (ABMP)



Algorithm 3 ABMP: Adaptive Blockwise Mixed Precision

- 1: Given block partition \mathcal{G} and importance \mathbf{Z}
 - 2: Compute $s_g = \sum_{(i,j) \in g} \mathbf{Z}_{ij}$ for all $g \in \mathcal{G}$
 - 3: Set $k \leftarrow \lfloor 0.05 |\mathcal{G}| \rfloor$ and sort blocks by s_g
 - 4: Assign $b_g = 3$ to top- k , $b_g = 1$ to bottom- k , others $b_g = 2$ ($\frac{1}{|\mathcal{G}|} \sum_g b_g = 2$)
 - 5: Apply DAQ quantization to each block using its assigned order $K = b_g$.
-



Q&A Datasets

Table 1: Results of GPTQ, GPTAQ, Slim-LLM, and our Quant-dLLM with 2-bit weight quantization among 7 tasks on LLaDA-Base, LLaDA-Instruct, LLaDA-1.5, Dream-Base, and Dream-Instruct. The numbers in parentheses represent the number used for evaluation. Best results are marked in **bold**.

Model	Method	MMLU(5)	Wino.(5)	PIQA(0)	ARC-C(0)	ARC-E(0)	Hella.(0)	BBH(0)	Avg
LLaDA-8B-Base	FP	65.76	74.59	74.70	44.11	74.20	54.27	42.60	61.46
	GPTQ	34.75	57.85	57.45	22.44	37.04	33.77	4.06	35.34
	GPTAQ	34.73	59.04	56.42	22.27	38.17	35.13	5.34	35.87
	Slim-LLM	47.98	60.85	62.46	26.11	53.41	39.25	6.64	42.39
	Quant-dLLM	56.87	68.19	69.75	36.26	68.69	46.45	32.18	54.06
LLaDA-8B-Instruct	FP	63.91	72.30	74.43	53.24	79.34	53.95	47.89	63.58
	GPTQ	42.15	58.88	61.21	27.22	51.18	37.46	3.93	40.29
	GPTAQ	44.94	60.22	61.04	29.01	52.44	37.18	4.55	41.34
	Slim-LLM	50.35	60.69	65.72	34.73	65.57	41.34	27.48	49.41
	Quant-dLLM	54.07	65.67	70.73	43.86	71.97	47.39	35.02	55.53
LLaDA-1.5	FP	64.21	54.38	74.54	54.27	79.80	54.38	56.27	62.55
	GPTQ	48.48	37.47	61.48	30.20	52.90	37.47	6.91	39.27
	GPTAQ	47.26	35.89	58.81	28.41	51.39	35.89	5.38	37.58
	Slim-LLM	44.99	41.86	65.89	33.70	64.86	41.86	24.90	45.44
	Quant-dLLM	54.32	47.82	71.16	47.18	73.06	47.82	38.09	54.21
Dream-7B-Base	FP	69.43	73.56	74.59	55.38	82.58	54.42	50.12	65.73
	GPTQ	23.84	50.51	53.86	19.88	28.70	29.69	3.08	29.94
	GPTAQ	23.90	52.49	55.88	20.82	30.43	30.98	4.00	31.21
	Slim-LLM	25.54	51.62	54.24	20.05	31.14	30.83	4.18	31.09
	Quant-dLLM	40.22	59.19	63.71	29.35	54.29	40.87	25.59	44.75
Dream-7B-Instruct	FP	69.71	72.93	75.14	57.25	83.80	54.43	57.92	67.31
	GPTQ	23.96	50.59	54.90	18.52	30.77	30.11	5.32	30.60
	GPTAQ	24.32	49.72	56.42	20.73	31.82	30.97	5.11	31.30
	Slim-LLM	25.29	51.93	56.09	20.56	36.99	31.55	7.61	32.86
	Quant-dLLM	43.14	57.93	66.59	33.62	62.96	41.36	30.35	47.99

Mathematical & Scientific Reasoning




Figure 3: Average accuracy of mathematical & scientific reasoning, and code generation datasets on LLaDA series and Dream series.



Ablation Study

- a) Effectiveness of our **Masked Calibration Simulation**.
- b) Effectiveness of **Adaptive Blockwise Mixed Precision**.
- c) Both **Data-aware Objective Reformulation** and **Row-column Successive Re-scaling** can improve the performance.
- d) Quant-dLLM remains stable across **varying calibration set sizes**.

Table 2: Ablation studies on LLaDA-8B-Base and Dream-7B-Base. We report MMLU in 5 shots.

(a) Effectiveness of MCS

Model	MCS	MMLU(5) ↑
LLaDA-8B-Base	✗	52.10
	✓	56.87
Dream-7B-Base	✗	37.81
	✓	40.22

(c) Ablation Study for DAQ

Model	DAQ Component	MMLU(5) ↑
LLaDA-8B-Base	baseline	39.26
	RSR w/o DOR	48.32
	RSR w/ DOR	56.87
Dream-7B-Base	baseline	27.84
	RSR w/o DOR	34.73
	RSR w/ DOR	40.22

(b) Effectiveness of ABMP

Model	ABMP	0%	5%	10%	15%
LLaDA-8B-Base	✗	54.32	–	–	–
	✓	–	56.87	55.87	53.01
Dream-7B-Base	✗	32.75	–	–	–
	✓	–	34.50	40.22	31.11

(d) Ablation Study for Calibration Set Size.

Model	Calib. Data Size	MMLU(5) ↑
LLaDA-8B-Base	64	54.49
	128	56.87
	256	55.59
Dream-7B-Base	64	38.24
	128	40.22
	256	39.64

More Result

Table 1: Effectiveness of MCS on previous methods. We report MMLU in 5 shots.

((a)) LLaDA-8B-Base

Model	MCS	Method	MMLU(5) \uparrow
LLaDA-8B-Base	\times	GPTQ	34.75
	\checkmark	GPTQ	36.75
	\times	GPTAQ	34.73
	\checkmark	GPTAQ	37.92
	\times	Slim-LLM	47.98
	\checkmark	Slim-LLM	50.11

((b)) Dream-7B-Base

Model	MCS	Method	MMLU(5) \uparrow
Dream-7B-Base	\times	GPTQ	23.84
	\checkmark	GPTQ	24.08
	\times	GPTAQ	23.90
	\checkmark	GPTAQ	24.68
	\times	Slim-LLM	25.54
	\checkmark	Slim-LLM	26.01

Memory

- **Compressed Memory**
 - Quant-dLLM achieves the smallest model size among all schemes, thanks to the memory-efficient row-column scaling factor design.

Table 3: Model size of LLaDA-8B-Base under different methods.

Model	Method	Bit	Memory (GB)
LLaDA-8B	FP16	16	16.09
	GPTQ	2	3.70
	Slim-LLM	2	3.72
	Quant-dLLM	2	3.69

Contribution

- We propose **Quant-dLLM**, a novel post-training extreme low-bit quantization framework for dLLMs.
 - We present **Masked Calibration Simulation**, reducing the distribution mismatch of calibration.
 - We introduce **Data-aware Any-order Quantizer**, including Data-aware Objective Reformulation (DOR) and Row-column Successive Re-scaling (RSR).
 - We propose **Adaptive Blockwise Mixed Precision**, concentrating representational capacity on the most critical model components.
- Quant-dLLM *outperforms SOTA* 2-bit PTQ methods while requiring *less memory*.

Thanks!