

Q-ROAR: Outlier-Aware Rescaling for RoPE Position Interpolation in Quantized Long-Context LLMs (Student Abstract)

Ye Qiao, Sitao Huang

University of California, Irvine, Department of EECS
Irvine, CA, 92697, USA
{yeq6, sitaoh}@uci.edu

Abstract

Extending LLM context windows is key for long-range tasks. RoPE-based position interpolation (PI) scales input length without retraining, and post-training quantization (PTQ) enables efficient deployment; however, combining PI with PTQ degrades accuracy due to long-context aliasing, dynamic-range dilation, axis-grid anisotropy, and outlier shifts that induce position-dependent logit noise. We give the first systematic analysis of PI+PTQ and propose two diagnostics: Interpolation Pressure (per-band phase-scaling sensitivity) and Tail Inflation Ratio (outlier shift from short to long contexts). We then introduce Q-ROAR, a RoPE-aware, weight-only stabilization that bands RoPE dimensions and lightly searches per-band scales for W_Q, W_K , with an optional symmetric variant. Q-ROAR needs only a tiny long-context dev set and no fine-tuning or kernel changes, recovering up to 0.7% accuracy and more than 14% GovReport perplexity reduction while preserving short-context performance.

Motivation and Problem

LLMs(Touvron 2023) increasingly rely on long contexts for summarization, RAG, code, chain-of-thought, and argentic workflows. Inference-time RoPE scaling (PI) (Peng et al. 2023; Qiao et al. 2025) extends the window is effective even without finetuning, but it is typically studied in full precision. Meanwhile, Post Training Quantization(PTQ)(Frantar et al. 2022; Lin et al. 2024; Wang et al. 2025) is essential for practical serving.

We observe that naively applying position interpolation (PI) to post-training quantized (PTQ) LLM models degrades accuracy both within and beyond the pretraining window (Figure 1). We compare quantized LLaMA-2-7B models with and without YaRN interpolation (Peng et al. 2023). In the non-interpolated control (Figure 1a), quantized models behave as expected, showing only modest degradation. In contrast, under YaRN, all quantized variants deteriorate and most sharply for the generic RTN (round-to-nearest) configuration. AWQ performs better out of the box than RTN, suggesting that explicit activation outlier handling is implicated in the robustness gap. Motivated by this, we conduct a principled analysis and attribute the failures to *coupling* between RoPE scaling and quantization: (i) *aliasing* as high-

frequency phases wrap; (ii) *dynamic-range dilation* that inflates pre-activation tails; (iii) *anisotropy* when axis-aligned quantizers operate on RoPE-rotated pairs; and (iv) *outlier shift/amplification*. Together, these effects induce position-dependent logit noise.

Interpolation Pressure and Tail Inflation

Most RoPE interpolation methods share a unified form:

$$\phi_i^{\text{scaled}}(m) = \omega_i \frac{f(m)}{s_i}, \quad s_i > 0, \quad (1)$$

where $f(\cdot)$ warps positions and s_i rescales per-dimension frequency. Let the training regime support $|m - n| \leq L_0$, target displacement D , and deviation $\varepsilon_i(D) = \omega_i \left(\frac{f(D)}{s_i} - D_0 \right)$. We define the sensitivity

$$\text{IP}_i = \left| \frac{\partial \varepsilon_i(D)}{\partial s_i} \right| = \omega_i \frac{f(D)}{s_i^2}, \quad (2)$$

which grows with ω_i and D , identifying high-frequency bands as fragile. To capture PI-induced outlier shift, we use *Tail-Inflation Ratios* at a high quantile $1 - \varepsilon$:

$$\text{TIR}_i^W = \frac{Q_{|w_i^\top h|, \text{long}}(1 - \varepsilon)}{Q_{|w_i^\top h|, \text{short}}(1 - \varepsilon)}, \quad \text{TIR}_i^A(m) = \frac{Q_{\|R(\phi_i^{\text{scaled}}(m))u_i\|_\infty}(1 - \varepsilon)}{Q_{\|R(\phi_i(m))u_i\|_\infty}(1 - \varepsilon)}. \quad (3)$$

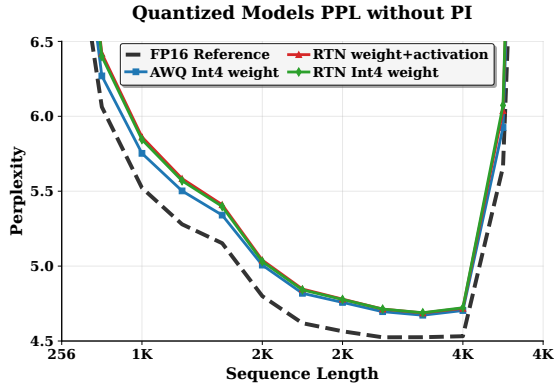
TIR^W reflects pre-activation tail growth; TIR^A reflects phase-axis amplitude inflation that increases activation clipping.

Geometric intuition. RoPE preserves ℓ_2 energy but rotates signal and quantization noise relative to axis-aligned bins; PI changes the angular trajectory over long spans. Thus, even uniform step sizes calibrated at short contexts become phase-misaligned at long contexts, increasing effective error along certain axes and perturbing attention logits.

Q-ROAR (Band-wise Weight Rescaling)

We partition RoPE pairs into B log-spaced frequency bands $\{\mathcal{B}_b\}_{b=1}^B$ and assign a single scale g_b :

$$W_Q^{(b)} \leftarrow g_b W_Q^{(b)}, \quad W_K^{(b)} \leftarrow \begin{cases} g_b W_K^{(b)} & \text{(shared mode)} \\ g_b^{-1} W_K^{(b)} & \text{(symmetric mode)} \end{cases} \quad (4)$$



(a) None Interpolated models

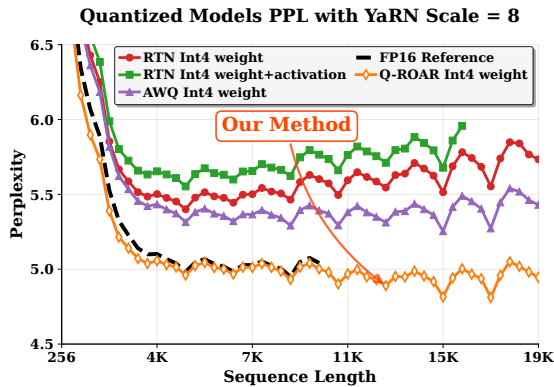
(b) YaRN interpolation ($s = 8$)

Figure 1: Quantized Llama-2-7b vs. GovReport PPL

Symmetric mode approximately preserves logit magnitude. We set per-band windows $\mathcal{G}_b = [g_b^{\min}, g_b^{\max}]$ by combining IP (tight for high frequencies) with TIR^W (shrink inflated tails). A tiny long-context dev set (~ 10 docs) drives a lightweight grid search that minimizes a length-weighted perplexity objective.

We choose **Band count** $B \in \{6, 8\}$ with log-frequency grouping; **Grid** 5–9 log-spaced candidates per band with coordinate and early stopping (gain $< \eta$); **Objective** $\sum_{L \in \mathcal{L}} w_L \text{ppl}(L; \{g_b\})$ with w_L emphasizing longer lengths; **Safety** bound g_b via $\gamma_b = 1 + \frac{\tau}{1 + \log(\omega_{b, \text{med}}/\omega_{\min})}$ and κ/TIR_b^W , with $\kappa \in [1.0, 1.3]$, $\tau \in [0.2, 0.5]$.

We focus on rescaling *weights* of key and query projection layers (W_Q, W_K) rather than adjusting activation for three reasons. (i) Activation statistics drift with content/position under PI; changing activation clips/steps couples to kernels/runtimes. (ii) Weight rescaling is static, quantizer and kernel agnostic (AWQ/RTN), and portable. (iii) Symmetric scaling keeps logit scale stable, avoiding retuning.

We sampled 10 long documents from Proof-pile (Zhangir Azerbayev 2022) dataset that each of them are longer than 60k tokens; cache short vs. PI distributions to compute TIR. **Complexity.** with running average sliding window size of 256. If each band has K candidates, evaluation cost is $\mathcal{O}(BK)$ runs (coordinate search) or $\mathcal{O}(K^B)$

Setting	Ctx	2K	4K	8K	16K	32K
Extended context with YaRN (64K tokens, $s=16$)						
FP16	64K	4.437	4.359	4.329	4.175	6.069
RTN W4	64K	4.544	4.485	4.470	4.485	6.713
AWQ W4	64K	4.489	4.421	4.405	4.414	6.302
Q-ROAR W4	64K	4.444	4.393	4.321	4.181	5.833

Table 1: GovReport perplexity on LLaMA-2-7B across evaluation context sizes (lower is better).

Setting	Ctx	5-shot Avg. (%)	WikiText-2 PPL
Base context (4096 window)			
Baseline FP16	4K	70.83	5.47
RTN W4	4K	64.14	5.73
RTN W4-A4	4K	62.02	5.92
AWQ W4	4K	64.25	5.61
AWQ W4-A4	4K	62.31	5.77
Extended context with YaRN (32K window, $s=8$)			
Baseline FP16	32K	63.71	6.09
RTN W4	32K	63.32	6.41
RTN W4-A4	32K	62.84	6.60
AWQ W4	32K	63.52	6.31
AWQ W4-A4	32K	62.53	6.49
Q-ROAR W4	32K	63.96	6.19
Q-ROAR W4-A4	32K	63.21	6.40

Table 2: LLaMA-2-7B on standard LLM benchmarks under different quantization and PI settings.

(small B only) with token reuse via sliding windows. **Serialization.** Store $\{g_b\}$ alongside checkpoint; apply at load time to (W_Q, W_K) tensors only. The search was conducted using two NVIDIA RTX 4090 and roughly took 4 GPU hours.

Experiments

For long-context evaluation, we benchmarked GovReport, and for standard LLM performance we used WikiText2 (Merity et al. 2016) and five zero-shot tasks. At the base 4K context, quantized models track FP16 with only modest degradation. Under YaRN interpolation (32K/64K), RTN and AWQ exhibit clear drops in accuracy and higher perplexity, confirming the destructive interaction between position interpolation and quantization. In contrast, Q-ROAR consistently improves robustness that recovering up to 0.7% accuracy on the standard suite and yielding over 14% relative perplexity reductions on GovReport versus RTN while preserving short-context performance. Overall, Q-ROAR mitigates aliasing and outlier amplification, enabling stable long context inference in quantized LLMs without retraining or kernel changes.

Conclusion

We show PI harms PTQ LLMs and introduce IP and TIR to measure RoPE–quantization coupling. **Q-ROAR** uses weight-only, bandwise rescaling of (W_Q, W_K) to deliver consistent long-context gains with negligible overhead and no kernel changes, mitigating aliasing and outlier shifts. It assumes RoPE and suits weight-quantized models. Q-ROAR cuts long-context perplexity by 14–21% while maintaining short-context accuracy and plug-in compatibility.

References

- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.
- Lin, J.; et al. 2024. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of machine learning and systems*, 6: 87–100.
- Merity, S.; et al. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Peng, B.; Quesnelle, J.; Fan, H.; and Shippole, E. 2023. Yarn: Efficient context window extension of large language models. *arXiv preprint arXiv:2309.00071*.
- Qiao, Y.; Xu, H.; Zhang, X.; and Huang, S. 2025. Rethinking RoPE Scaling in Quantized LLM: Theory, Outlier, and Channel-Band Analysis with Weight Rescaling. *arXiv preprint arXiv:2510.00028*.
- Touvron, H. o. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, Y.; Qiao, Y.; Huang, S.; and Kwon, H. 2025. APEX-Q: Arbitrary-dimension Product-EXTension Quantization for Accelerated LLM Deployment. *Authorea Preprints*.
- Zhangir Azerbayev, B. P., Edward Ayers. 2022. Proof-pile.