Multilinguality after Quantization: Low-Bits, Calibration Languages, Mixed-Precisions, and Cross-Linguality

Anonymous ACL submission

Abstract

There are inconsistent conclusions about the impact of calibration data on Post-Training Quantization (PTQ) across studies, specifically regarding the factors that make the method effective, such as language that we explore in this work. We investigate this question by evaluating GPTQ and its mixed-precision variants across three low-bit settings, using language-targeted calibration sets across five language families. By evaluating performance across both perplexity and downstream tasks, we find that selecting an appropriate calibration language is generally beneficial, though model sensitivity varies with bit-width, target language, and task difficulty. We further show that mixed-precision techniques expand the representational capacity of low-bit models, occasionally allowing 2-bit models to approach 3-bit performance when paired with language-targeted calibration. Furthermore, we explore the correlation between these factors and observe that the calibration set sensitivity is influenced by different combinations of factors for each controlled variable, which can be explained by the downstream task results. Finally, while the benefits of language-targeted calibration transfer to sibling languages within the same family, the magnitude of this transfer is not reliably predicted by closeness in family alone. Together, these findings show calibration data selection as a deliberate design choice in multilingual PTQ rather than an incidental one.

1 Introduction

The impressive multilingual capabilities of Large Language Models (LLMs) scale with model size (Lin et al., 2022; Yang et al., 2025; Dang et al., 2024; Grattafiori et al.); as a result, deploying the model becomes expensive. To address this, Post-Training Quantization (PTQ) has become a widely used compression technique, as it offers a highly efficient compression method while retaining the original model's performance by conditioning on a small calibration dataset (e.g., 128 samples) (Gholami et al., 2021; Rakka et al., 2025). However, existing works primarily focus on English performance, thus making the field of quantization on low-resource languages underexplored, especially at extreme quantization levels (3- and 2-bit) (Marchisio et al., 2024; Zeng et al., 2025; Kurz et al., 2025; Choenni and Titov, 2025). Furthermore, since PTQ is a non-training approach, optimizing the model for a specific target language is challenging, leaving the single-pass calibration set as the primary adjustable hyperparameter.

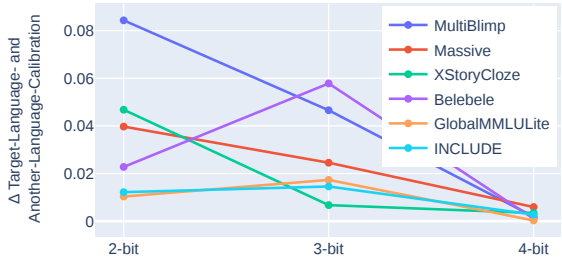

Figure 1: When evaluating on target language, Quantization using mono-target-language calibration set generally outperforms mono-another-language, while the sensitivity varies between bit-width and task difficulty.

Despite the intuitive potential to improve PTQ via calibration set composition, empirical studies report varying results across problem sets. Interestingly, the perceived impact of calibration sets has shifted chronologically, from negligible to substantial results. Rather than dismissing these diverse outcomes as random noise or coincidence, the underlying causes of these discrepancies demand closer analysis.

In this work, we offer a new perspective on the discrepant effects of calibration set by investigating multiple variables. We want to answer the following question: *Which previously uneval-*

uated variables can systematically predict how a calibration set affects a quantized model's performance? Specifically, we experiment using `GPTQ` (Frantar et al., 2023a) as the representative `PTQ` method, evaluating it across three low bit-levels (4-bit, 3-bit, and 2-bit). To smooth the transition between bit-levels, we incorporate both "light" and "heavy" mixed-precision variants (§3.1) (Huang et al., 2025; Xiao et al., 2025). Finally, we evaluate the effectiveness of language-targeted calibration across five language families, focusing on the relative improvement over other-language calibration.

Our primary contributions are as follows:

- We show that the effect of language-specific calibration is generally favorable relative to other-language calibrations. The advantage is clear in perplexity, although it is rarely substantial in downstream tasks.
- We observed that mixed-precision methods help stabilize the model's sensitivity to the calibration set. In some cases, they enable lower-bit models to compete with higher-bit counterparts, particularly when combined with language-targeted calibration.
- We investigate four parameters that modulate the sensitivity of language-specific calibration: bit-levels, downstream task difficulty, mixed-precision variants, and target languages. Our findings suggest that each controlled variable has a variable combination to increase the impact of calibration set.
- We examine the benefits of language-specific calibration transfer to sibling languages within the same family. Generally, the effect is amplified for lower-resource languages.

To the best of our knowledge, only a few studies address this specific problem. Most notably, concurrent works by Kurz et al. (2025) and Chimoto et al. (2026) reach similar conclusion although in a different compression setting (§2).

2 Related Works

Quantization. Quantization is an optimization problem that seeks to preserve performance degradation despite the reduced representational capacity, motivating the development of several quantization algorithms (Gholami et al., 2021). These methods are broadly categorized into three types: Data-Free Quantization (e.g., BitsandBytes (Dettmers et al., 2022)), Post-Training Quantization (PTQ) (e.g., `GPTQ` (Frantar et al., 2023a) and AWQ (Lin et al., 2026)), and Quantization-Aware Training (QAT) (e.g., ternary-bit methods (Ma et al., 2024) and `EfficientQAT` (Chen et al., 2025)). Most PTQ and QAT methods are guided by calibration data, with QAT further requiring substantial retraining to recover performance loss. Additionally, some PTQ have succeeded in achieving extremely low-bit quantization, such as `BRECQ` (Li et al., 2021) and `QuIP` (Chee et al., 2024; Tseng et al., 2024). To further mitigate performance degradation, mixed-precision methods assign higher bit-widths to important weights (Rakka et al., 2025).

Effect of calibration set selection on PTQ. While most quantization research focuses on algorithmic improvements, a growing body of literature begins to analyze the impact of the calibration dataset on the quantized model. Early studies suggests that the effect of the calibration set is negligible in general use cases (Paglieri et al., 2024; Williams and Aletras, 2024). We note that these works focus on 4-bit quantization, which is widely regarded as the bit that works in retaining performance (Dettmers and Zettlemoyer, 2023; Jin et al., 2024). However, it remains unclear whether calibration set can actively mitigate the performance drop of lower-bit models.

Recent studies have incorporated extreme-bits in their evaluations, revealing a much more pronounce effect of task-targeted calibration set selection. Liu et al. (2024) demonstrating substantial fluctuations when calibrating on tasks and their subjects, while not consistently produce optimal result. Conversely, Liu et al. (2025a) demonstrate substantial improvements on heavy reasoning task, particularly when using `GPTQ`.

Calibration Sets and Multilinguality. While the impact of compression on multilingual LLMs has been the subject of ongoing study (Marchisio et al., 2024), the specific role of the calibration set has only recently been explored. For example, Zeng et al. (2025), Choenni and Titov (2025), and Chimoto et al. (2026) propose multilingual calibration sets to better preserve general multilingual capabilities.

Subsequent studies focus on language-targeted compression using calibration datasets. Kurz

et al. (2025) investigate language-specific calibration through pruning, a compression technique that removes redundant model parameters. Similar to our PTQ setting, Chimoto et al. (2026) investigate how varying the language composition of calibration data affects outlier activation statistics, focusing on 4-bit precision and vanilla PTQ methods. Fundamentally, these studies arrive at a shared conclusion: the language distribution within the calibration set improves multilingual perplexity, while downstream performance gains are not guaranteed. Our work distinguishes itself through its analytical focus, broader quantization configurations, and preliminary cross-lingual analysis.

3 Experimental Setup

3.1 Quantization Techniques

We primarily focus on PTQ because we observe the impact of language in the calibration set. Specifically, we use `GPTQ` as a representative PTQ method, along with two mixed-precision variants that differ in how aggressively they protect important weights, enabling smooth transitions across bit-widths. The "light" variant (`SliM-LLM`, (Huang et al., 2025)) assigns important weights a higher bit-width than the rest, but still below 16 bits. The "heavy" variant (`TaCQ`, (Xiao et al., 2025)) goes further and keeps important weights in full FP16. Detailed discussions of each method are deferred to Appendix A.

3.2 Model, Calibration Dataset, and Languages

Model. We use open-source multilingual LLMs that have 8B parameter variants, namely `Qwen3 8B` (Yang et al., 2025), `Llama 3.1 Instruct 8B` (Grattafiori et al.), and `Aya Expanse 8B`[1] (Dang et al., 2024). We then quantize each model into 3 bit-levels: 4-bit, 3-bit, and 2-bit. Although a 2-bit model is currently impractical, we seek to explore the relationships between variables to better understand how calibration impacts PTQ.

Calibration Dataset and Languages. We use a machine translation dataset, Flores Plus `dev` set (NLLB Team et al., 2024), as the source for our quantization calibration data. We do not enforce exact semantic alignment across calibration sets (except through random seed), as each language has varying token length for the same sentence, while calibration set has a sequence length requirement. We select English (EN), Tamil (TA), Indonesian (ID), Swahili (SW), and Chinese (ZH), which span diverse language families, with Tamil and Chinese using non-Latin scripts.

Evaluation Tools. We use `lm-eval-harness` (Gao et al., 2024) for downstream tasks, with `vLLM` (Kwon et al., 2023) as the inference backend for higher throughput. Token perplexity is computed using the Hugging Face evaluate library (Von Werra et al., 2022). We report the geometric mean across datasets, since perplexity values can become extremely large in the 2-bit setting, and the arithmetic mean would be dominated by these outliers.

Hyperparameter. Our hyperparameter, calibration set construction setup, and quantization failure can be seen in the Appendix C.

3.3 Multilingual Tasks

Perplexity. We use Flores Plus `devtest` set (NLLB Team et al., 2024) to evaluate the multilingual perplexity of the quantized models. We use this data since it provides a fair comparative multilingual performance due to its parallel structure.

Multilingual Downstream Tasks. We evaluate on various task difficulties to explore the language-specific calibration effectiveness across configurations. We evaluate on minimal grammatical pair test with MultiBlimp and its language variants (Jumelet et al., 2025; Liu et al., 2025b; Leong et al., 2023; Kryvosheieva and Levy, 2025), scenario classification with MASSIVE (FitzGerald et al., 2022; Bastianelli et al., 2020), story completion with XStoryCloze (Lin et al., 2021). Furthermore, we evaluate on knowledge-explicit reading comprehension with Belebele (Bandarkar et al., 2024), and both knowledge-implicit question answering with Global MMLU Lite (Singh et al., 2025) and INCLUDE[2] (Romanou et al., 2024). All evaluations are conducted in a multiple-choice format without using thinking mode for any of models used.

[1] Although `Aya` primarily supports English, Indonesian, and Chinese (while `Llama` officially supports only English from the selections), we found that both downstream task performance remains reasonably strong across the evaluated languages.

[2] We prioritize evaluation on the 5 languages, while some task may not covering these languages. Specifically for INCLUDE, we use aggregated german and dutch as the english evaluation, as they are from the Germanic language family.

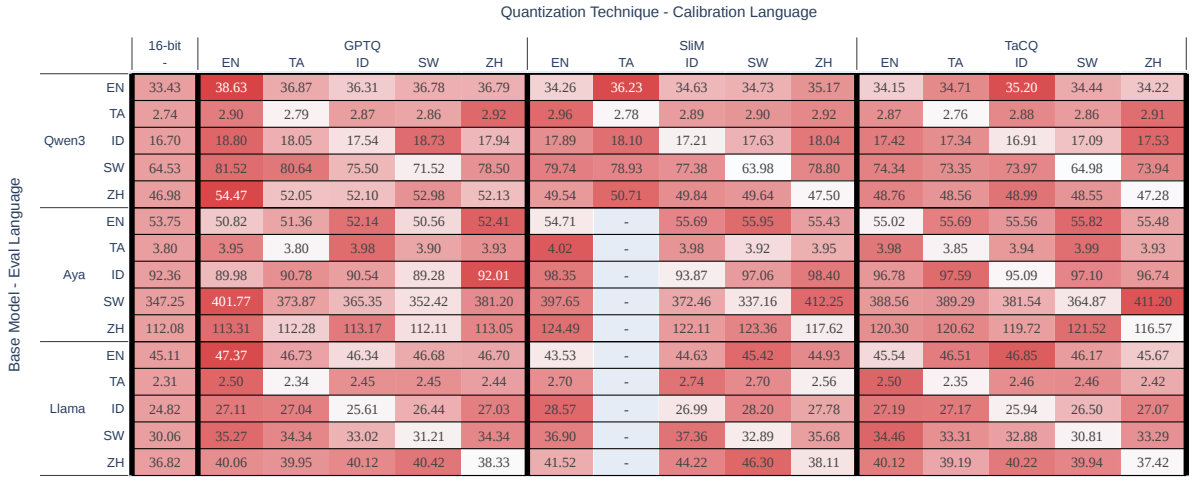

Figure 2: Multilingual Perplexity on 4-bit models. The local-diagonal cells represent the relative effect of Target-Calibration over Cross-Calibrations, i.e., a single row of each quantization method (lighter/lower is better). The relative advantage is clear even in the 4-bit settings, while we defer the rest lower-bits to Appendix D.1. `SliM LLM` `Llama` Tamil produce outlier, thus biasing the perplexity result.

3.4 Paper's Shorthand

For brevity, we introduce the following shorthand notation. C_{X} (X Calibration Language): An X-only calibration set. C_{target}, C_{cross}, or C_{default}: On evaluating a target language (X), using the same (C_{X}), different ($C_{\neg X}$), or default (C_{en}) calibration set. Consequently, $Target$, $Cross$, and $Default$: the impact of corresponding calibration set. `Unquantized`: The original 16-bit model.

4 Results and Analysis

4.1 The effect of Direct-Calibration on Perplexity

Effects of C_{target} are Reflected in Multilingual Perplexity. Consistent with prior work by Kurz et al. (2025) and Chimoto et al. (2026), our results for `4-bit` models show that the C_{target} perplexity over unquantized baseline is small, typically within a 10% difference. However, C_{target} is generally better than C_{cross} across configurations. This pattern is clear on the diagonal entries of Figure 2. However, some exception do occurs, such as `GPTQ` on English. We defer the extreme-bit evaluation heatmaps in the Appendix D.1.

Effect of C_{target} becomes significant at extreme-bit widths. We compare the perplexity reductions when switching from C_{default} to C_{target} in Table 1. In general, the improvement made by C_{target} increases as bit-width decreases, amplify by the decreasing performance of C_{default}, although the magnitude varies across languages. Notably, Indonesian is the least sensitive to calibration language, while Swahili is the most sensitive, despite both languages using Latin script. Furthermore, sensitivity increases sharply when moving from 3-bit to 2-bit (e.g., `Qwen3 GPTQ` on Indonesian rises from 21.1% to 96.8%), which is far more severe than from 4-bit to 3-bit (from 6.7% to 21.1%). This contrast is more pronounced for non-Latin scripts (e.g., `SliM-LLM` on Tamil and Chinese). These results

Model	Quant	Bit	Perplexity Reduction TA	ID	SW	ZH
Qwen3 8B	GPTQ	4	3.9%	6.7%	12.3%	4.3%
		3	32.6%	21.1%	50.1%	10.8%
		2	99.98%	96.8%	96.6%	43.0%
	SliM	4	6.1%	3.8%	19.8%	4.1%
		3	42.2%	28.8%	55.7%	13.2%
		2	99.94%	91.4%	89.4%	99.97%
	TaCQ	4	3.9%	2.9%	12.6%	3.0%
		3	27.6%	16.9%	43.6%	9.8%
		2	78.9%	81.8%	82.4%	56.6%
Aya Expanse 8B	GPTQ	4	3.8%	−0.6%	12.3%	0.2%
		3	28.2%	17.8%	44.9%	27.1%
		2	99.98%	92.1%	92.8%	95.2%
	SliM	4	—	4.6%	15.2%	5.5%
		3	—	17.01%	46.4%	16.98%
		2	—	99.7%	94.7%	99.996%
	TaCQ	4	3.3%	1.7%	6.1%	3.1%
		3	14.9%	7.0%	28.1%	10.0%
		2	73.7%	62.8%	77.0%	61.3%

Table 1: Perplexity reduction when adopting C_{target} across quantization settings (higher reduction is less robust towards calibration). Calculated by $(C_{\text{en}} - C_{\text{target}})/C_{\text{en}}$. The green and red colored denote the least and most sensitive language towards C_{target}.

Prompt: இலையுதிர் காலத்தின் நடுவில் குடும்பத்துடன் சுற்றுலா செல்வது பற்றி ஒரு கவிதை எழுதுங்கள்

Write me a poem about the picnic with family in the middle of the fall season

---

`GPTQ` 2-bit **English**-calibrated (48256.105)

矜toString Mid representations furthermore rematch convo阶 '\''riel整个_foot whereneau响应矩 ...

`GPTQ` 2-bit **Tamil**-calibrated (8.060)

உயம், எனக் ஆக்ராஞ்சாடி வர்டின் ஒருங்கரூடி காரியை பயன்பாடு காக்காசேன் என்று ஒரு பெரிய ச❓

Uyam, I am a big fan of Agrachandi Ward's Orungarudi Karikai Kakkasen.

`TaCQ` 2-bit **English**-calibrated (14.578)

Okay, let's break down this sentence. First, the user is asking about the process of writing a sentence in a way that it's not a sentence, but ...

`TaCQ` 2-bit **Tamil**-calibrated (3.077)

கவித்துருவை நிறுத்துவது: இலையுதிர் காலத்தின் நடுவில், குடும்பத்துடன் சுற்றுலா செல்லும் சிறு வி❓

Stopping the Poet: In the middle of autumn, A little girl goes on a family trip.

Table 2: Contrasting the generation of Tamil when the model is calibrated using C_{en} and C_{ta} on `Qwen3 8B`. Number inside parenthesis is the model's Tamil perplexity. Higher-bit models adhere to the prompt.

suggest a strong degradation in target-language linguistic capability under C_{default} at 2-bit quantization, particularly for non-Latin scripts. If this linguistic capability collapse can be mitigated, 2-bit performance may become comparable to 3-bit models.

Calibration Language and Mixed-Precision determine 2-bit model generation. To highlight this effect, we report generated Tamil sentences, which shows the most contrast between `GPTQ` and `TaCQ` (see Table 2). The outputs vary depending on whether Target-Calibration, Mixed-Precision, or both are used. In particular, C_{target} steers generation towards the target language, while mixed-precision improves overall coherence. Interestingly, the `TaCQ` 2-bit C_{en} model generates reasoning in the default language (i.e., English) but falls into repetitive looping behavior. This may indicate that C_{default} preserves reasoning ability but degrades broader capabilities, such as long-context comprehension and duplicate word penalty. Consequently, using both Mixed-Precision and C_{target} allows the 2-bit model to function as a targeted monolingual model, albeit generating overly simplistic sentences.

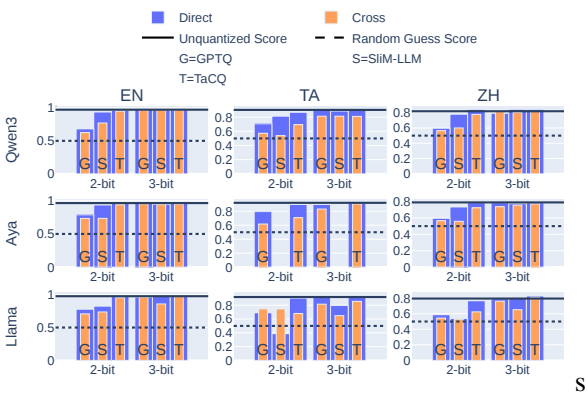

Figure 3: MultiBlimp: Contrasting accuracy (Higher is better) of $Target$ vs mean of 2 $Cross$. Notably, 2-bit can be comparable to 3-bit when use both mixed-precision and target-calibration. We omit the Swahili (SW) and Indonesian (ID) as their unquantized accuracy are near random guess (0.5).

4.2 The effect of Target-Calibration on downstream task performance

We observe that the effectiveness of C_{target} varies across task difficulties; it is critical for some while negligible for most. This suggests a fluctuating calibration set sensitivity across quantized configurations. We present results comparing C_{target} and C_{cross} as bar-charts.

2-bit models require mixed-precision and Target-Calibration on `MultiBlimp`. Figure 3 shows the MultiBlimp evaluation results (see Appendix D.2 for Massive and XStoryCloze), revealing three relationships between C_{target}, languages resource level, and mixed-precisions. Notice that `3-bit` models achieve near-lossless performance, while 2-bit models exhibit high variance, suggesting that MultiBlimp is a suitable benchmark for 2-bit quantization evaluation. First, 2-bit models can match 3-bit performance when using techniques such as mixed-precision and Target-Calibration; mixed-precision is primary for recovery, while C_{target} is secondary for performance gains (e.g., `Qwen3` on ZH). Second, the behaviors of mixed-precisions and C_{target} differ between language resource level: Low-resource language requires C_{target} for stability, hence mixed-

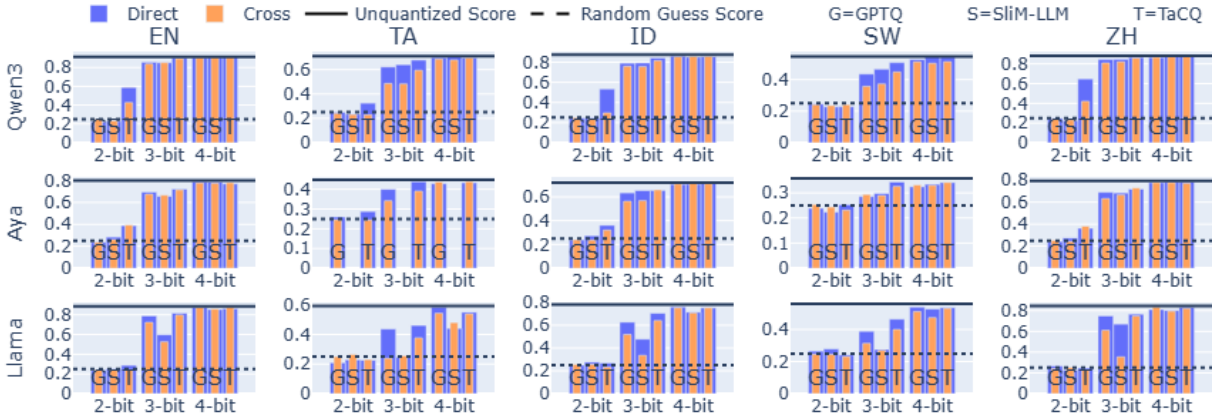

Figure 4: Belebele: Contrasting accuracy (Higher is better) of C_{target} vs mean of 4 C_{cross}. Overall, the relative advantage remains, while absolute improvements mostly occur on 3-bit and low-resource language. Notably, English is robust towards calibration language.

precision is not enough on TA; while, C_{target} remained secondary in High-resource languages (e.g., contrast `TaCQ` on TA and ZH). Third, the effect of light and heavy mixed-precisions are noticeably different. These three findings serve as foundational patterns that persist in subsequent, more difficult tasks. Finally, the delta performance between C_{target} and C_{cross} ranges from -0.05 to 0.27.

3-bit models sometimes show strong calibration sensitivity on `Belebele`, especially in low-resource languages. `Belebele` is a suitable benchmark for evaluating 3-bit models, due to the less informative result on 4-bit and 2-bit, as shown in Figure 4 (see Appendix D.3 for INCLUDE and Global MMLU Lite). The three patterns identified in the previous tasks are also observed here, although reduced in intensity. First, mixed-precision remains primary to be beneficial, as observed in the `TaCQ` results across 3-bit configurations. Furthermore, C_{target} improves performance retention for low-resource languages (e.g., `Qwen3` on Tamil) and acts as a safeguard against unexpected quantization failure (e.g., `Llama` with `SliM-LLM` across all languages). Next, light and heavy mixed-precisions remain noticeably different. Interestingly, 2-bit models can achieve decent performance when paired with `TaCQ` and Target-Calibration, although the outcome depends on the base model. This pattern suggests that stronger base model capabilities may enhance resilience to extreme quantization. Overall, the delta performance at `3-bit` models ranges from -0.02 to 0.32.

Suitable tests for 4-bit configurations. Existing benchmarks fail to significantly differentiate the C_{target} and C_{cross} effects at `4-bit` models. Such evaluation likely require advanced reasoning or language generation skills, such as mathematics or code generation. This speculation is drawn from a recent study by Liu et al. (2025a), which shows that task-calibration effects are pronounced in heavy reasoning tasks, particularly when using `GPTQ`. However, those substantial improvements limited to 3-bit models and English task.

4.3 The effect of Target-Calibration on within-family perplexity

We extend our study to cross-lingual intra-family scenarios, serving as a preliminary test of whether C_{target} preserves performance across similar, related languages. We analyze the effect on perplexity using all languages available in Flores Plus `test` set (NLLB Team et al., 2024), leaving exhaustive downstream tasks evaluations for future work (see preliminary in Appendix F.1).

We limit our in-depth analysis to the `Austronesian` (Indonesian) language family. We investigate the correlation between the degree of transfer and its relationship to the target language based on sibling relationship (i.e. glottolog[3] (Hammarström et al., 2025).) and geographical proximity.

Austronesian (Indonesian) cross-lingual mainly

[3]https://glottolog.org/

Figure 5: Austronesian (Indonesian) cross-lingual transfer benefits aggregated across quantization methods. Values represent the perplexity ratio percentage between C_{id} and mean of 4 C_{cross}. Language are ordered (left to right) according to their phylogenetic proximity in Glottolog. Among the evaluated languages, only 5 languages from Indonesian (ID) to Acehnese (ACE) that are closer in siblingness, while the remaining languages meet at the highest node.

applies to languages spoken in Indonesia, Malaysia, and Philippines. The Austronesian language family (Indonesian) includes seven phylogenetic subgroups, ranging from the highest level Austronesian node to the terminal node, Standard Indonesian. Almost all available Malay and Indonesian languages share a common ancestral subgroup within the second-highest (`Malayo-Polynesian`) branch, except for four Indonesian languages: Balinese (BAN), Buginese (BUG), Javanese (JAV), and Sundanese (SUN). Following our naming convention, the 13 languages meet at the `Malayo-Polynesian` level as distant-relative ($ID_{\text{gene,distant}}$), and the 5 more phylogenetically closer languages as close-relative ($ID_{gene,close}$).

From the 2-bit section in Figure 5, the cross-lingual transfer can be categorized into four levels: High, Moderate, Low, and Negligible. Specifically, we classify it from the transfer ratio percentage: 0-20% (C_{cross} is $5\times$ higher than C_{target}), 20-50%, 50-80%, and 80+%. These results suggest that transfer levels cannot be predicted solely by sibling or geographical proximity. For instance, while Indonesian and Malay exhibit a high transfer effect, not all language spoken in Indonesia or all Inner-Languages automatically exhibit a moderate or high transfer effect. Specifically, Acehnese (ACE) ($ID_{\text{gene,close}}$) exhibits a low transfer effect, together with Buginese (BUG) ($ID_{\text{gene,distant}}$). Conversely, Javanese (JAV) and Sundanese (SUN) ($ID_{\text{gene,distant}}$) exhibit a moderate transfer effect. Geograhically, the former two originated outside Java Island, where the Indonesia's capital is located, while the latter two originated on the island itself. Interestingly, four languages spoken in Philippines are categorized as having a low transfer effect, while the remaining languages (Oceanic languages and Malagasy) exhibit a negligible transfer effect.

Variable	Order	Comment
Model's release data	Newer	Bias towards the newest
Quantization Technique	Heavier	See subsection 3.1
Bit Width	Higher	Omit 4-bit
Language Resource	Higher	See Table 4 in Appendix
Task Difficulty	Harder	Based on similar results
Delta Acc	Higher	$Target - Cross$

Table 3: Ordinal Setting for spearman correlation.

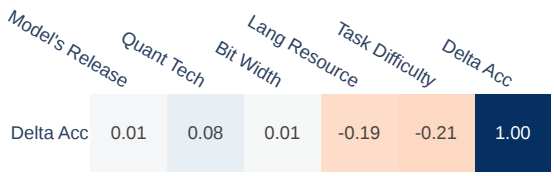

Figure 6: Which variable influence calibration sensitivity?

Other Language Family analysis.. We defer the rest intra-family transfer in Appendix F.1. In summary, we observe that high-resource languages tend to restrict benefits to themselves, yielding negligible or even negative effects for their sibling languages. Conversely, mid and low-resource languages, especially tamil that use different scripts, exhibit much stronger cross-lingual benefits. We note the limitation that these cross-lingual effects are observed only through perplexity and negligible except at 2-bit level.

4.4 Calibration Sensitivity Analysis

To uncover factors that influence model's calibration sensitivity, we calculate spearman correlation among variables[4][5] as detailed in Table 3. We omit the correlation between variable settings, as they only serve as an experimental co-occurrence. Please take notice that this correlation does not directly translate to "improving the quantization accuracy" but a general calibration sensitivity, i.e., $C_{\text{target}} - C_{\text{cross}}$. For example, later we see that negative correlation in quantization technique does not mean using `TaCQ` reduce the accuracy, but interpreted as `TaCQ` is more robust than `GPTQ`, as shown in the downstream task results.

General Correlation. Figure 6 shows the overall spearman correlation. We can see that only task difficulty and language resource (weakly negatively) influence the delta accuracy, although bit width and quantization technique should also influence the delta, as seen in the downstream task results. Thus, we need to analyze isolated variable.

Grouping by Bit. Bit Width is the simplest perspective that one can take, for example "If I want to quantize to N-bit, when should I care for language-targeted calibration". Figure 7 shows 2-bit and 3-bit correlation. At 2-bit, easier task and heavier quantization technique will make Target-Calibration effective. Conversely, it does not matter what evaluation language one wants to evaluate at 2-bit. As for 3-bit, only lower-resource languages are needed to use Target-Calibration, while weak negative correlation in quantization technique can have a meaning of mixed-precision methods are more robust than GPTQ. These patterns match with downstream task evaluations.

Grouping by others variables. There are some unexplained interactions among other variables that we defer to Appendix E. Overall, there is not a single variable that influence calibration sensitivity in all configurations. Notably, Calibration sensitivity in English is weak-negative towards harder task and higher bits, which most calibration set study falls in this category.

[4] 4-bit bias towards negative correlation.

[5] To simplify ranking order on Task Difficulty, we use our own judgement from our empirical results and reading the dataset, i.e., MultiBlimp, Massive, and XStoryCloze as 1, while Belebele, INCLUDE, and Global MMLU Lite as 2.

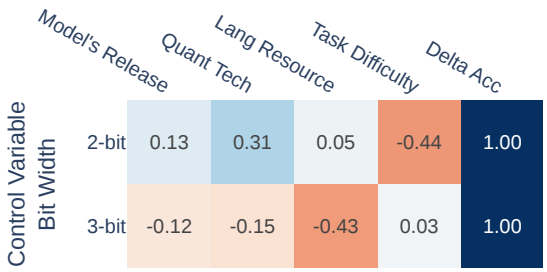

Figure 7: Calibration sensitivity across Bit Widths. Each row represents the correlation between delta accuracy and another variable.

5 Conclusion

Our work demonstrates that models quantized via `PTQ` *are affected* by the choice of language in the calibration set, leading to several key insights. First, we recommend using Target-Calibration for monolingual LLM applications. However, the improvements over Default- or Cross-Calibration are often marginal, while notable exceptions exist. Second, we find that a quantized model's sensitivity to the calibration set is primarily limited by its representational capacity. Specifically, mixed-precision methods expand this representational capacity, occasionally allowing 2-bit models to perform comparably to their 3-bit counterparts, while Target-Calibration maximizes performance within that capacity. Third, the effects of Target-Calibration do transfer across languages within the same language family (intra-family), although the degree depends on language resource level. Fourth, the correlation scores support the downstream task findings. Finally, our findings suggest that the conflicting calibration impacts in previous studies may be explained by the tendency of English calibration sensitivity to decrease as bit-level and task difficulty increase.

Suggestion

We observe that assessing the true target-language capability of quantized LLMs remains challenging. Therefore, we encourage the NLP community to develop dedicated language capability benchmarks, analogous to CEFR language proficiency levels, that can quantitatively measure the linguistic competence of smaller models, including quantized LLMs. Such benchmarks may be more feasible to construct given that existing online language learning materials have been categorized according to CEFR proficiency levels.

Limitation

We limit our study to three 8B models, thus the observed patterns may not generalize well to larger models. Second, although we use five languages spanning varying resource levels and language families, our selection is not exhaustive. In particular, non-Latin script languages warrant further investigation. Third, the distinction between simple tasks and more complex reasoning tasks is based on a qualitative inspection of dataset instances, although a more rigorous quantitative characterization of each task difficulty may be possible. Fourth, although recent studies suggest that GPTQ family exhibits higher sensitivity to calibration set, we have yet to evaluate whether this observation generalizes across a broader family of quantization methods. In particular, future work should examine it holds under more recent PTQ methods that have achieved strong performance on extreme-bit settings. Fifth, while Floresplus is a translation benchmark, we have not conducted an ablation study on the effect of semantic parallelism. Lastly, the cross-lingual analysis is based on geographic and sibling similarities, while more direct measures of linguistic similarity have yet to be explored.

Ethical Considerations

LLM Usage: We use Gemini and ChatGPT to proofreading and correcting: grammar, word choice, and awkward phrasing.

Potential Risk: As we found that language-targeted calibration on 2-bit TaCQ can generate coherent response in target language, albeit simple, we afraid that it can cause unintended multilingual future risk, such as multilingual jailbreak (Deng et al., 2024). With assumption that the model can provide knowledgeable answer in extreme-bit setting on top of target language understanding.

References

Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2024. The belebele benchmark: a parallel reading comprehension dataset in 122 language variants. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 749–775, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.

Emanuele Bastianelli, Andrea Vanzo, Pawel Swietojanski, and Verena Rieser. 2020. SLURP: A spoken language understanding resource package. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7252–7262, Online. Association for Computational Linguistics.

Manuel Brack, Malte Ostendorff, Pedro Ortiz Suarez, José Javier Saiz, Iñaki Lacunza Castilla, Jorge Palomar-Giner, Alexander Shvets, Patrick Schramowski, Georg Rehm, Marta Villegas, and Kristian Kersting. 2024. Community OSCAR: A community effort for multilingual web data. In *Proceedings of the Fourth Workshop on Multilingual Representation Learning (MRL 2024)*, pages 232–235, Miami, Florida, USA. Association for Computational Linguistics.

Jerry Chee, Yaohui Cai, Volodymyr Kuleshov, and Christopher De Sa. 2024. Quip: 2-bit quantization of large language models with guarantees. *Preprint*, arXiv:2307.13304.

Mengzhao Chen, Wenqi Shao, Peng Xu, Jiahao Wang, Peng Gao, Kaipeng Zhang, and Ping Luo. 2025. Efficientqat: Efficient quantization-aware training for large language models. *Preprint*, arXiv:2407.11062.

Everlyn Asiko Chimoto, Mostafa Elhoushi, and Bruce A. Bassett. 2026. Calibrating beyond english: Language diversity for better quantized multilingual llm. *Preprint*, arXiv:2601.18306.

Rochelle Choenni and Ivan Titov. 2025. M-wanda: Improving one-shot pruning for multilingual llms. *Preprint*, arXiv:2505.21171.

John Dang, Shivalika Singh, Daniel D'souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, and 26 others. 2024. Aya expanse: Combining research breakthroughs for a new multilingual frontier. *Preprint*, arXiv:2412.04261.

Yue Deng, Wenxuan Zhang, Sinno Jialin Pan, and Lidong Bing. 2024. Multilingual jailbreak challenges in large language models. *Preprint*, arXiv:2310.06474.

Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. Llm.int8(): 8-bit matrix multiplication for transformers at scale. *Preprint*, arXiv:2208.07339.

Tim Dettmers and Luke Zettlemoyer. 2023. The case for 4-bit precision: k-bit inference scaling laws. *Preprint*, arXiv:2212.09720.

Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha

Britan, Wouter Leeuwis, Gokhan Tur, and Prem Natarajan. 2022. Massive: A 1m-example multilingual natural language understanding dataset with 51 typologically-diverse languages. *Preprint*, arXiv:2204.08582.

Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. 2023a. Gptq: Accurate post-training quantization for generative pre-trained transformers. *Preprint*, arXiv:2210.17323.

Elias Frantar, Sidak Pal Singh, and Dan Alistarh. 2023b. Optimal brain compression: A framework for accurate post-training quantization and pruning. *Preprint*, arXiv:2208.11580.

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac'h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024. The language model evaluation harness.

Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W. Mahoney, and Kurt Keutzer. 2021. A survey of quantization methods for efficient neural network inference. *Preprint*, arXiv:2103.13630.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. The llama 3 herd of models.

Harald Hammarström, Robert Forkel, Martin Haspelmath, and Sebastian Bank. 2025. glottolog/glottolog: Glottolog database 5.2.1.

Wei Huang, Haotong Qin, Yangdong Liu, Yawei Li, Qinshuo Liu, Xianglong Liu, Luca Benini, Michele Magno, Shiming Zhang, and Xiaojuan Qi. 2025. Slim-llm: Salience-driven mixed-precision quantization for large language models. *Preprint*, arXiv:2405.14917.

Renren Jin, Jiangcun Du, Wuwei Huang, Wei Liu, Jian Luan, Bin Wang, and Deyi Xiong. 2024. A comprehensive evaluation of quantization strategies for large language models. *Preprint*, arXiv:2402.16775.

Jaap Jumelet, Leonie Weissweiler, Joakim Nivre, and Arianna Bisazza. 2025. Multiblimp 1.0: A massively multilingual benchmark of linguistic minimal pairs. *Preprint*, arXiv:2504.02768.

Dhruva Kashyap, Chaitanya Murti, Pranav K Nayak, Tanay Narshana, and Chiranjib Bhattacharyya. 2026. Modhifi: Identifying high fidelity predictive components for model modification. *Preprint*, arXiv:2511.19566.

Daria Kryvosheieva and Roger Levy. 2025. Controlled evaluation of syntactic knowledge in multilingual language models. In *Proceedings of the First Workshop on Language Models for Low-Resource Languages*, pages 402–413, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Simon Kurz, Jian-Jia Chen, Lucie Flek, and Zhixue Zhao. 2025. On the limitations of language targeted pruning: Investigating the calibration language impact in multilingual llm pruning. *Preprint*, arXiv:2408.14398.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Wei Qi Leong, Jian Gang Ngui, Yosephine Susanto, Hamsawardhini Rengarajan, Kengatharaiyer Sarveswaran, and William Chandra Tjhi. 2023. Bhasa: A holistic southeast asian linguistic and cultural evaluation suite for large language models. *Preprint*, arXiv:2309.06085.

Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. 2021. Brecq: Pushing the limit of post-training quantization by block reconstruction. *Preprint*, arXiv:2102.05426.

Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2026. Awq: Activation-aware weight quantization for llm compression and acceleration. *Preprint*, arXiv:2306.00978.

Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O'Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona Diab, and 2 others. 2022. Few-shot learning with multilingual language models. *Preprint*, arXiv:2112.10668.

Xi Victoria Lin, Todor Mihaylov, Mikel Artetxe, Tianlu Wang, Shuohui Chen, Daniel Simig, Myle Ott, Naman Goyal, Shruti Bhosale, Jingfei Du, Ramakanth Pasunuru, Sam Shleifer, Punit Singh Koura, Vishrav Chaudhary, Brian O'Horo, Jeff Wang, Luke Zettlemoyer, Zornitsa Kozareva, Mona T. Diab, and 2 others. 2021. Few-shot learning with multilingual language models. *CoRR*, abs/2112.10668.

Ruikang Liu, Yuxuan Sun, Manyi Zhang, Haoli Bai, Xianzhi Yu, Tiezheng Yu, Chun Yuan, and Lu Hou. 2025a. Quantization hurts reasoning? an empirical study on quantized reasoning models. *Preprint*, arXiv:2504.04823.

Yijun Liu, Yuan Meng, Fang Wu, Shenhao Peng, Hang Yao, Chaoyu Guan, Chen Tang, Xinzhu Ma, Zhi Wang, and Wenwu Zhu. 2024. Evaluating the generalization ability of quantized llms: Benchmark, analysis, and toolbox. *Preprint*, arXiv:2406.12928.

Yikang Liu, Yeting Shen, Hongao Zhu, Lilong Xu, Zhiheng Qian, Siyuan Song, Kejia Zhang, Jialong Tang, Pei Zhang, Baosong Yang, Rui Wang, and Hai Hu. 2025b. A systematic assessment of language models with linguistic minimal pairs in chinese. *Preprint*, arXiv:2411.06096.

Shuming Ma, Hongyu Wang, Lingxiao Ma, Lei Wang, Wenhui Wang, Shaohan Huang, Li Dong, Ruiping Wang, Jilong Xue, and Furu Wei. 2024. The era of 1-bit llms: All large language models are in 1.58 bits. *Preprint*, arXiv:2402.17764.

Kelly Marchisio, Saurabh Dash, Hongyu Chen, Dennis Aumiller, Ahmet Üstün, Sara Hooker, and Sebastian Ruder. 2024. How does quantization affect multilingual LLMs? In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 15928–15947, Miami, Florida, USA. Association for Computational Linguistics.

ModelCloud.ai and qubitium@modelcloud.ai. 2024. Gpt-qmodel. `https://github.com/modelcloud/gptqmodel`. Contact: qubitium@modelcloud.ai.

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, and 20 others. 2024. Scaling neural machine translation to 200 languages. *Nature*, 630(8018):841–846.

Davide Paglieri, Saurabh Dash, Tim Rocktäschel, and Jack Parker-Holder. 2024. Outliers and calibration sets have diminishing effect on quantization of modern llms. *Preprint*, arXiv:2405.20835.

Irina Proskurina, Luc Brun, Guillaume Metzler, and Julien Velcin. 2024. When quantization affects confidence of large language models? *Preprint*, arXiv:2405.00632.

Mariam Rakka, Marios Fournarakis, Olga Krestinskaya, Jinane Bazzi, Khaled N. Salama, Fadi Kurdahi, Ahmed M. Eltawil, and Mohammed E. Fouda. 2025. Mixed-precision quantization for language models: Techniques and prospects. *Preprint*, arXiv:2510.16805.

Angelika Romanou, Negar Foroutan, Anna Sotnikova, Zeming Chen, Sree Harsha Nelaturu, Shivalika Singh, Rishabh Maheshwary, Micol Altomare, Mohamed A. Haggag, Snegha A, Alfonso Amayuelas, Azril Hafizi Amirudin, Viraat Aryabumi, Danylo Boiko, Michael Chang, Jenny Chim, Gal Cohen, Aditya Kumar Dalmia, Abraham Diress, and 40 others. 2024. Include: Evaluating multilingual language understanding with regional knowledge. *Preprint*, arXiv:2411.19799.

Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David I. Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Wei-Yin Ko, Sebastian Ruder, Madeline Smith, Antoine Bosselut, Alice Oh, Andre F. T. Martins, Leshem Choshen, and 5 others. 2025. Global mmlu: Understanding and addressing cultural and linguistic biases in multilingual evaluation. *Preprint*, arXiv:2412.03304.

Shivalika Singh, Freddie Vargus, Daniel Dsouza, Börje F. Karlsson, Abinaya Mahendiran, Wei-Yin Ko, Herumb Shandilya, Jay Patel, Deividas Mataciunas, Laura OMahony, Mike Zhang, Ramith Hettiarachchi, Joseph Wilson, Marina Machado, Luisa Souza Moura, Dominik Krzemiński, Hakimeh Fadaei, Irem Ergün, Ifeoma Okoh, and 14 others. 2024. Aya dataset: An open-access collection for multilingual instruction tuning. *Preprint*, arXiv:2402.06619.

Albert Tseng, Jerry Chee, Qingyao Sun, Volodymyr Kuleshov, and Christopher De Sa. 2024. Quip#: Even better llm quantization with hadamard incoherence and lattice codebooks. *Preprint*, arXiv:2402.04396.

Leandro Von Werra, Lewis Tunstall, Abhishek Thakur, Sasha Luccioni, Tristan Thrush, Aleksandra Piktus, Felix Marty, Nazneen Rajani, Victor Mustar, and Helen Ngo. 2022. Evaluate & evaluation on the hub: Better best practices for data and model measurements. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 128–136, Abu Dhabi, UAE. Association for Computational Linguistics.

Qianli Wang, Mingyang Wang, Nils Feldhus, Simon Ostermann, Yuan Cao, Hinrich Schütze, Sebastian Möller, and Vera Schmitt. 2026. Through a compressed lens: Investigating the impact of quantization on factual knowledge recall. *Preprint*, arXiv:2505.13963.

Miles Williams and Nikolaos Aletras. 2024. On the impact of calibration data in post-training quantization and pruning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 10100–10118. Association for Computational Linguistics.

Hanqi Xiao, Yi-Lin Sung, Elias Stengel-Eskin, and Mohit Bansal. 2025. Task-circuit quantization: Leveraging knowledge localization and interpretability for compression. *Preprint*, arXiv:2504.07389.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41

others. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Hongchuan Zeng, Hongshen Xu, Lu Chen, and Kai Yu. 2025. Multilingual brain surgeon: Large language models can be compressed leaving no language behind. *Preprint*, arXiv:2404.04748.

A Background

A.1 Quantization

Model quantization aims to reduce memory and computational requirements while preserving the original performance. It seems to be a simple problem but a challenging one in practice. Doing this naively inevitably alters the weight distribution and therefore propagates the error to the final layer, leading to performance degradation. The existing methods range from simple lookup table, intricate mathematical optimizations, to resource-intensive retraining (Gholami et al., 2021; Rakka et al., 2025).

Despite their differences, quantization fundamentally seeks an optimal function that maps a model's parameters from a high-precision to a low-precision format while ensuring the mapped model maximally emulates the original behavior. Moreover, due to the inherent compromises of this trade-off, theoretically, quantization can be optimized to maintain general capabilities (which may result in a shared global performance loss) or to preserve a *niche* capability (which currently requires deliberate configuration). For example, one may prioritize preserving the model's broad multilingual proficiency or its proficiency in a single targeted-language. While most quantization methods focus on the former, for our problem, the latter is preferable.

A.2 Post-Training Quantization

Utilizing a calibration set allows one to optimize the direction of quantization process; most notably, it enables the simulation of the quantized model's behavior. The primary metric to quantify the error during quantization is the layer-wise reconstruction error, defined in Equation A.2. The fundamental objective is to minimize the layer-wise output distance between the unquantized (original) model and its quantized counterpart, in the hope of achieving an effective yet efficient quantization process. Recent finding has proven that layer-wise optimization is linearly upper-bounded (at maximum linear multiplication) of the end result (Kashyap et al., 2026).

$$\ell(W, \hat{W}, l) = ||X_l W_l - X_l \widehat{W_l}||_2$$

The symbols in the equation are as follows: X_l represents the l-th layer's input batch, while W_l and $\widehat{W_l}$ represent the unquantized and quantized l-th layer's weight tensors, respectively. $||\cdot||_p$ denotes p-norm or, in simple term, a distance metric.

`GPTQ` (Frantar et al., 2023a) and its origin, `OBQ` (Frantar et al., 2023b), navigate layer-wise reconstruction optimization problem using the mathematical properties of Hessian matrix. In this context, the Hessian matrix (second-order derivatives) holds the sensitivity of each parameter to induce error when changed (e.g., during quantization). Specifically, `GPTQ` quantizes each parameter iteratively and optimizes the remaining unquantized weights through the inverse Hessian matrix, which informs *inverse* sensitivity to balance the backlash of point-quantization. **More importantly, each layer's Hessian matrix is calculated solely from the layer's input matrix, which intuitively made the calibration set important**. Although the original work states that `GPTQ` is robust to calibration set sampling, recent discoveries empirically prove otherwise.

A.3 Improvements: Mixed-Precision

Notice that, from the Hessian matrix, we can estimate which weights induce the most sensitive to error than the rest, implying that not every weight is equally important. It is an observed phenomena, such as outlier-inducing weights (Dettmers et al., 2022). Therefore, we can mix the precisions in at least two ways: maintaining the total bit count (such as `SliM-LLM` (Huang et al., 2025)) or quantizing all weights except the most important ones (such as `TaCQ` (Xiao et al., 2025)). For convenience, we call the former as Light Mixed Precision and the latter as Heavy Mixed Precision. Fundamentally, both variants introduce a weight-importance determination method step prior to quantization.

`SliM-LLM`, Salience-Driven Mixed-Precision Quantization (Huang et al., 2025) focuses on two aspects: structured mixed-precision and preserving the total bit count. First, structured mixed-precision ensures that the model can be inferred directly in low-precision through the currently supported efficient hardware operations, which necessitate weights having the same precision for each input dimensions (cite). Second, by splitting $2D$ weights into groups of consecutive input dimensions (e.g., each group consisting of 128 parameters of output dimension), `SliM-LLM` quantizes the top-k most important groups to $N+1$ bits, the least-k important groups to $N-1$

bits, and the rest to the target bit, or N. The salience of each group is determined by averaging the group's Hessian matrix. We observe that `SliM-LLM` is prone to overfitting on calibration set due to the exhaustive nature of determining the k parameter.

`TaCQ`, Task-Circuit Quantization (Xiao et al., 2025) is inspired by Mechanistic Interpretability methods. Specifically, `TaCQ` uses absolute weight, the difference between unquantized (clean) and quantized (corrupted) weights, and calibration data's gradient to determine important weights, as expressed in Equation A.3. Moreover, the authors reported that task-specific calibration set improves related task accuracies by saving task-related weights in higher-precision. We observed that `TaCQ` is more robust than the other two `GPTQ` variants, probably due to its heavy mixed-precision and the use of gradient salience. However, the original `TaCQ` method does not yet support inference directly in low-bit due to its unstructured mixed-precision.

$$Salience = W \odot |W - \widehat{W}| \odot G$$

Here, W is the unquantized weight matrix (clean), $\widehat{W}$ is the vanilla `GPTQ` quantized weight matrix (corrupted), G is the gradient of chosen loss function, and $\odot$ represents the Hadamard (item-wise) multiplication operation.

B Expanded Related Works

B.1 Quantization Degrades Quality Beyond Downstream Metrics

Numerous studies suggest that PTQ can achieve "negligible performance loss" on standard benchmarks (Frantar et al., 2023a; Dettmers and Zettlemoyer, 2023; Jin et al., 2024). However, recent studies indicates that this claim does not hold true across more nuanced qualitative attributes. For instance, Marchisio et al. (2024) observe that while quantized models exhibit negligible performance degradation on automated generative tasks, they receive lower scores from human evaluators, particularly when generating non-Latin scripts. Wang et al. (2026) employ interpretability tools to show that the factual memorization abilities of quantized LLMs degrade modestly. Furthermore, Proskurina et al. (2024) highlight the shift in confidence score, resulting in overconfidence on incorrect answers. Finally, Liu et al. (2025a) find that extreme quantization substantially degrades the reasoning capability of 3-bit models.

C Experiment Setup

Hyperparameters Generally, we adopt the default parameters and scripts from the respective GitHub repositories: `GPTQModel`[6](ModelCloud.ai and qubitium@modelcloud.ai, 2024) (v4.2.0), `SliM-LLM`[7] (Aug 2024), and `TaCQ`[8] (June 2025). We set the random seed to `1234`, the number of calibration samples to `128`, and the sequence length to `2048`. Additionally, for `TaCQ`, we use `16-bit float datatypes` during backpropagation to mitigate GPU memory constraints. Unfortunately, both mixed-precision methods fail to pack the quantized models into lower-bit formats (e.g., saving weight natively in 2-bit); thus, we must simulate the quantized models in full-precision during inference. Notably, we quantize and evaluate each configuration in a single run with a fixed random seed, cutting both budget and runtime.

Calibration Set Construction We follow the methodologies of `SliM-LLM` and `TaCQ` to construct the calibration set. For each language, we first concatenate all dataset instances, separated by a double newline character (`\n\n`). Then, the concatenated text is tokenized using the corresponding model's tokenizer. Finally, we randomly select chunks of 2048 consecutive tokens as individual samples.

Library Modification: We locally extend the SliM-LLM and TaCQ libraries with additional model families and utilities for the FloresPlus calibration dataset.

Quantization Failure `SliM-LLM Aya` experience numerical instability error, although we have tried multiple dampening coefficient.

Language and Average Bits Figure 4 shows our work's language coverage. The resource level is taken from `Aya dataset` (Singh et al., 2024), while the order is taken from OSCAR[9] (Brack et al., 2024). Figure 5 shows the average bit of each quantization techniques. Notably, even though `TaCQ` allocates some weight into 16-bit, its average bit stay close to the `GPTQ` counterpart (less than 0.1 point).

[6]https://github.com/ModelCloud/GPTQModel
[7]https://github.com/Aaronhuang-778/SliM-LLM
[8]https://github.com/The-Inscrutable-X/TACQ
[9]https://huggingface.co/datasets/oscar-corpus/community-oscar

Language	Characteristic		Ranking	
	Script	Family	Order	Resource Level
English	Latin	Germanic	5	High
Chinese	Hans	Sino-Tibetan	4	High
Indonesian	Latin	Austronesian	3	Mid
Tamil	Tamil	Dravidian	2	Mid
Swahili	Latin	Niger-Congo	1	Low

Table 4: Language resource ranking

Target Bit	Avg bit (Higher Bit Count)		
	GPTQ	SliM	TaCQ
2-bit	2.0	2.0 (5.52%)	2,049 (0.35%)
3-bit	3.0	3.0 (5.21%)	3,046 (0.35%)
4-bit	4.0	4.0 (25.05%)	4,042 (0.35%)

Table 5: Average bits of each quantization techniques and target-bit levels. `SliM-LLM` maintains total-bits while optimizing its bit-allocation, allocating higher bit to $target + 1$. `TaCQ` keep its higher-bits in 16-bit precision. Total parameter = 6.9B

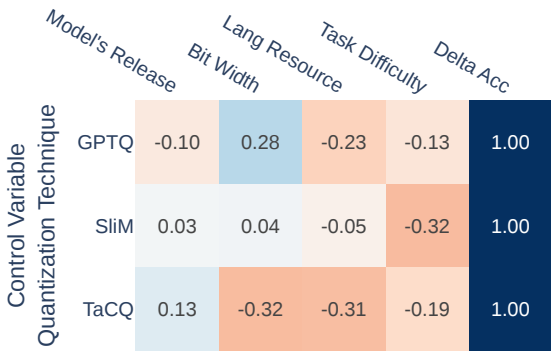

Figure 8: Calibration sensitivity across Quantization Techniques. Each row represents the correlation between delta accuracy and another variable.

D Expanded Results

D.1 Perplexity Results

Figure 13 shows perplexity results at 3-bit and 2-bit. The C_{target} advantage over C_{cross} persist in extreme-bit, even with a higher contrast.

D.2 Massive and XStoryCloze Task Results

We can observe that the substantial effects of C_{target} are still visible in Massive (shown in Figure 14) and less visible in XStoryCloze (shown in Figure 15). This can be attributed to the hypothesis that calibration set can have substantial effect when the potential of quantized model is present, i.e., the gap between unquantized and the quantized. While in XStoryCloze results, we can see that the gap between Unquantized score and random-guess threshold are close (<0.2 accuracy point on ID, SW, and ZH). Another way to explain it, the original score is close to the random-guess, therefore there is nothing left to calibrate.

D.3 Global MMLU Lite and INCLUDE Task Results

Figure 16 and Figure 17 show Global MMLU Lite and INCLUDE results, respectively. We observe that the C_{target} effect are visible in some configurations, in particular at 3-bit, similar to Belebele evaluation in Figure 4. The combination of mixed-precision and Target-Calibration can serve as a safeguard against extreme-bit quantization failure, such in `3-bit SliM-LLM` model on INCLUDE's ID. Another note is that German and Dutch evaluation at 3-bit in INCLUDE prefer Cross-Calibration rather than using English-Calibration. This can be attributed to the fact that English-Calibration has less effective transfers on Germanic languages as seen in Figure 18.

E Extended Calibration Sensitivity Analysis

E.1 Grouping by Quantization Technique

Figure 8 show GPTQ, SliMLLM, and TaCQ correlation settings. Bit width has an interesting influence on quantization techniques: Calibration sensitive towards GPTQ in higher bits while TaCQ in lower bits (Bit width correlation is positively weak in GPTQ, none in SliM, while negatively weak in TaCQ). Next, GPTQ and TaCQ weakly favor lower-resource, while only SliMLLM favor easier tasks. Respecting the MultiBlimp result (Figure 3), SliM-LLM is senstive towards calibration set across language, in particular on easier tasks.

E.2 Grouping by Language Resource

Figure 9 show the correlations viewed on each languages. The least fluctuative correlation, Indonesian and Chinese do not favor any variables toward calibration sensitivity. This corresponds to the downstream tasks results, as both of them less sensitive towards language-targeted calibration across variables. As for the most sensitive towards calibration language, Swahili and Tamil favor higher-bit width and easier task for calibration sensitivity. In particular for Tamil, it favors Qwen3 model, as though it can comes from a bias on our experiment that we conduct more evaluations on `Qwen3` than `Llama`. Lastly, English has an inverted sensitivity on higher bit-width and harder task, preferring

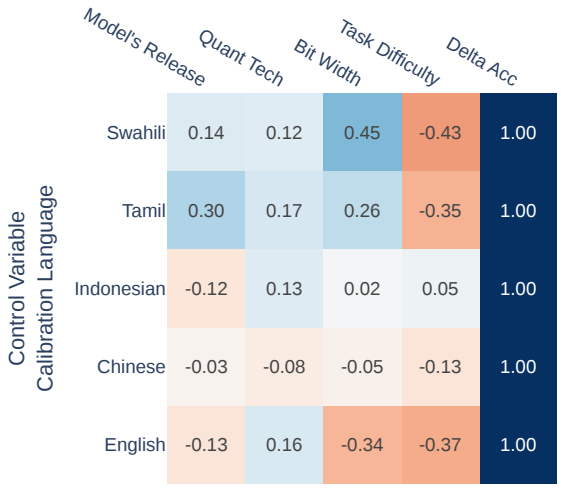

Figure 9: Calibration sensitivity across Calibration Languages. Each row represents the correlation between delta accuracy and another variable.

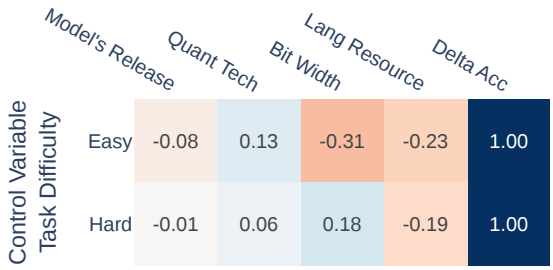

Figure 10: Calibration sensitivity across Task Difficulties. Each row represents the correlation between delta accuracy and another variable.

C_{cross} rather than C_{en}, such can be seen in 3-bit on Global MMLU Lite and INCLUDE.

E.3 Grouping by Task Difficulty

Figure 10 shows the Easier and Harder task correlation. Easier tasks undoubtly favor lower-bit width and lower-resource, plainly seen in task figures. However, it favor very weakly on quantization technique, although we've seen substantial delta in SliM-LLM and TaCQ while negligible in GPTQ across easy tasks. It means that the data find the easier task's sensitivity leaning more in bit-width than quantization techniques. As for the harder task correlation, it shows milder sensitiveness towards target-calibration.

F Expanded Cross-Lingual within family

Sino-Tibetan (Chinese) cross-lingual transfer may produce negative effect Based on the patterns in Figure 11, we categorize the cross-lingual effects into three levels: High, Low, and Detrimental (observed only at 2-bit setting). The high transfer effect are Chinese (ZH) and Yue Chinese (YUE), the two languages are the only Inner-Languages. Moderate effects are observed for Tibetan (BOD) and Dzongkha (DZO). The remaining languages fall into the detrimental effect, where a specific cross-lingual transfers yield significantly better perplexity. Upon closer inspection, all four in the Detrimental effect category tend to prefer calibration using a low-resource language: Swahili yields better perplexity for Southern Jinghpaw (KAC) and Mizo (LUS), while Tamil yields better perplexity for Manipuri (MNI) and Burmese (MYA). Interestingly, all four of these languages are geographically close to Myanmar (Burmese). Notably, these suboptimal effects are observe only at the 2-bit level. This may suggest a nuanced role for low-resource languages in improving another low-resource languages.

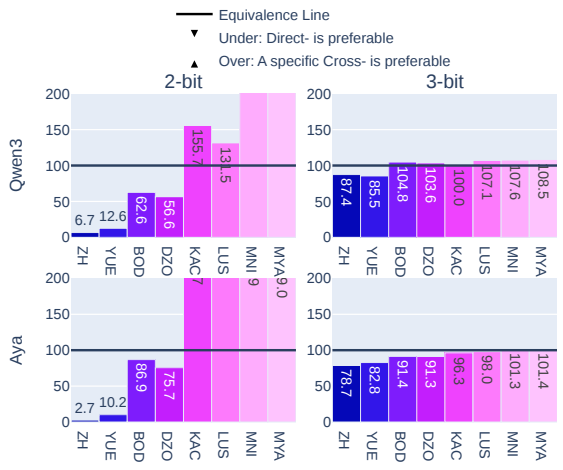

Figure 11: Sino-Tibetan (Chinese) cross-lingual transfer benefits aggregated across quantization methods. Values represent the perplexity ratio percentage between C_{zh} and mean of 4 C_{cross}. Languages are ordered (left to right) according to phylogenetic proximity in `Glottolog`. Among the evaluated languages, only 2 languages, ZH and Yue Chinese (YUE), that are closer in siblingness.

Cross-lingual effects are negligible or negative for high-resource languages, limited for Indonesian, and significant for low-resource languages. Across these five language families, we can see a general pattern (see Appendix F.1 for the barplot). High-resource languages tend to restrict benefits to themselves, yielding negligible or even

negative effects for their sibling languages. Conversely, lower-resource languages, especially those using different scripts, exhibit much stronger cross-lingual benefits. We note the limitation that these cross-lingual effects are observed only through perplexity remain and negligible except at 2-bit level; therefore, these findings do not override our general recommendation to use C_{target} for targeted language settings.

F.1 Cross-Lingual Perplexity Results

The pattern of high-resource languages tend to prefer transfer to self-language that can be seen in Figure 18, showing Germanic evaluations. We can see that at 2-bit, English is the only language that benefited from English-Calibration, even Eastern Yiddish (YDD) prefer Tamil language. The pattern of low-resources tend to share the benefit across language within the same family can be seen in Figure 12 and Figure 19, showing Dravidian (Tamil) and Niger-Congo (Swahili), respectively.

F.2 Cross-Lingual Downstream Task Results

Figure 20-23 show the general (barplot) and detailed (heatmap) evaluations of Cross-Lingual Calibration on Belebele, encompassing Austronesian (Indonesian) and Sino-Tibetan (Chinese). Due to the time constraint, we only evaluate on 3-bit for Austronesian and Qwen3 at three bit-levels for Sino-Tibetan. Overall, the transfer effect is negligible, clearly seen from the barplots, while the relative advantage varies according to the language family. Specifically, Austronesian still prefer Indonesian-Calibration (Swahili as secondary). As for Sino-Tibetan, evaluating on BOD and KAC have their unquantized's accuracy close to random-guess, so their evaluation should not be considered. As for Burmese (MYA), it prefers Tamil (Swahili as Secondary) for the extreme-bits. This pattern corroborates with the perplexity evaluations at Figure 11, although its belebele's delta accuracies are negligible.

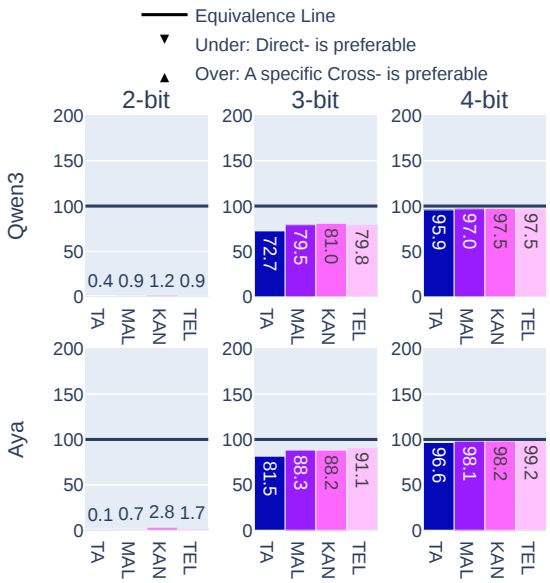

Figure 12: Perplexity of Dravidian (Tamil) on cross-lingual calibration. The score is calculated using $C_{\text{target}}/\overline{C_{\text{cross}}}$ (lower is preferring Tamil), aggregate across quantization methods.

Quantization Technique - Calibration Language

Base Model - Eval Language		16-bit	GPTQ					SliM					TaCQ				
		-	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH
Qwen3	EN	33.43	40.19	42.01	42.39	44.79	40.87	38.00	44.63	40.98	45.26	44.25	34.71	38.88	36.95	37.24	36.94
	TA	2.74	4.33	2.92	3.92	3.77	4.13	5.02	2.90	3.99	4.00	4.24	3.85	2.79	3.45	3.28	3.66
	ID	16.70	24.30	22.62	19.17	22.73	24.87	26.60	23.04	18.94	22.94	26.92	21.07	19.33	17.50	19.19	21.26
	SW	64.53	160.42	144.26	143.36	79.99	151.04	168.99	135.08	135.64	74.92	168.43	117.20	106.38	102.48	66.06	117.45
	ZH	46.98	62.90	62.48	62.97	66.25	56.10	62.77	62.12	63.99	68.86	54.45	53.50	54.06	54.49	53.38	48.26
Aya	EN	53.75	55.72	60.00	58.33	56.88	56.49	70.33	-	70.16	71.45	71.28	56.60	59.78	57.99	59.71	58.37
	TA	3.80	5.49	3.94	4.94	4.90	5.05	5.19	-	4.85	4.71	4.96	4.60	3.91	4.43	4.39	4.44
	ID	92.36	122.90	170.29	100.95	109.50	108.38	128.24	-	106.43	117.15	123.70	104.22	107.60	96.90	105.49	106.00
	SW	347.25	701.94	866.68	601.33	386.51	628.57	694.41	-	581.60	372.34	698.01	488.42	494.79	505.83	351.03	561.06
	ZH	112.08	170.11	262.47	151.64	153.63	124.03	153.81	-	148.35	152.57	127.69	132.58	135.09	132.25	139.73	119.32
Llama	EN	45.11	54.00	56.98	56.64	55.01	56.13	51.66	-	52.61	82.71	55.08	48.27	50.73	50.05	51.24	50.03
	TA	2.31	4.62	2.48	3.81	3.84	3.67	6.09	-	4.95	8.44	3.89	4.27	2.43	3.46	3.55	3.12
	ID	24.82	49.03	42.82	32.11	38.43	45.47	80.62	-	33.44	237.45	52.09	39.77	35.06	28.62	34.72	35.46
	SW	30.06	76.30	63.39	54.52	36.94	70.29	140.05	-	64.68	127.10	79.51	56.28	46.88	43.46	34.77	50.87
	ZH	36.82	65.91	57.60	62.21	62.64	45.39	107.70	-	68.90	493.79	44.50	57.11	48.48	52.05	56.76	41.28

Quantization Technique - Calibration Language

Base Model - Eval Language		16-bit	GPTQ					SliM					TaCQ				
		-	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH
Qwen3	EN	33.43	1.3e+04	4.5e+04	3.4e+04	4.3e+04	3.2e+04	136.42	9.5e+03	1.9e+03	1.3e+03	4.1e+03	44.43	63.76	51.82	65.68	60.22
	TA	2.74	4.8e+04	8.06	2.7e+04	3.8e+04	2.1e+04	7.9e+03	4.30	2.8e+04	1.3e+03	6.2e+03	14.58	3.08	11.10	9.83	11.66
	ID	16.70	4.0e+04	4.9e+04	1.3e+03	4.0e+04	3.4e+04	2.5e+03	8.5e+03	213.79	734.79	1.8e+04	118.65	92.52	21.58	44.40	135.16
	SW	64.53	3.8e+04	-	3.1e+04	1.3e+03	3.7e+04	7.2e+03	6.6e+03	4.3e+03	755.83	1.4e+04	391.55	636.81	233.38	68.74	351.53
	ZH	46.98	3.4e+04	5.2e+04	4.5e+04	-	2.0e+04	1.0e+06	9.2e+04	6.0e+05	2.0e+05	296.79	143.96	118.76	120.67	189.92	62.54
Aya	EN	53.75	235.53	831.15	525.80	724.30	908.81	126.67	-	242.23	405.97	9.6e+03	69.12	91.79	81.00	97.78	87.57
	TA	3.80	2.9e+04	6.78	5.1e+03	1.9e+03	3.8e+03	7.4e+05	-	5.7e+05	1.8e+06	6.6e+05	15.33	4.03	11.75	10.50	11.09
	ID	92.36	7.1e+03	6.5e+03	563.14	4.2e+03	5.4e+03	6.5e+04	-	222.24	1.3e+03	5.5e+05	305.45	213.56	113.63	204.26	287.21
	SW	347.25	2.1e+04	2.2e+04	1.6e+04	1.5e+03	1.8e+04	8.4e+03	-	5.0e+03	440.48	8.8e+04	1.5e+03	1.5e+03	1.3e+03	342.51	2.2e+03
	ZH	112.08	2.1e+04	1.1e+04	1.0e+04	1.3e+04	1.0e+03	6.7e+06	-	2.2e+06	1.4e+06	263.67	352.03	268.35	299.08	410.46	136.31
Llama	EN	45.11	1.3e+03	8.5e+03	6.1e+03	1.0e+04	7.7e+03	1.3e+03	6.5e+03	2.5e+04	4.7e+04	7.0e+03	71.31	123.02	94.74	108.40	110.84
	TA	2.31	-	9.53	-	-	5.7e+04	1.3e+07	5.14	3.0e+06	5.9e+05	2.1e+05	98.91	2.79	11.29	15.90	8.78
	ID	24.82	3.2e+04	3.9e+04	2.8e+03	2.7e+04	3.6e+04	9.0e+04	2.9e+04	8.4e+04	1.3e+05	5.5e+04	409.74	254.87	48.68	135.91	481.67
	SW	30.06	3.4e+04	4.6e+04	1.9e+04	4.1e+03	4.3e+04	4.8e+04	1.5e+04	3.4e+04	6.8e+04	2.0e+04	422.78	329.59	208.37	43.17	469.15
	ZH	36.82	3.9e+04	3.2e+04	3.7e+04	4.7e+04	3.1e+03	1.2e+07	6.7e+04	1.9e+06	3.9e+06	2.4e+03	698.26	276.42	381.74	1.2e+03	68.79

Figure 13: Floresplus perplexity: 3-bit (Top) and 2-bit (Bottom), lighter or lower is better.

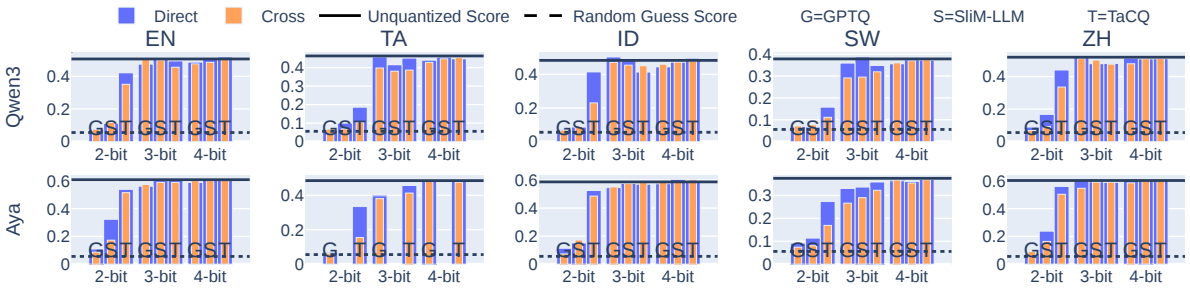

Figure 14: Massive Barplot

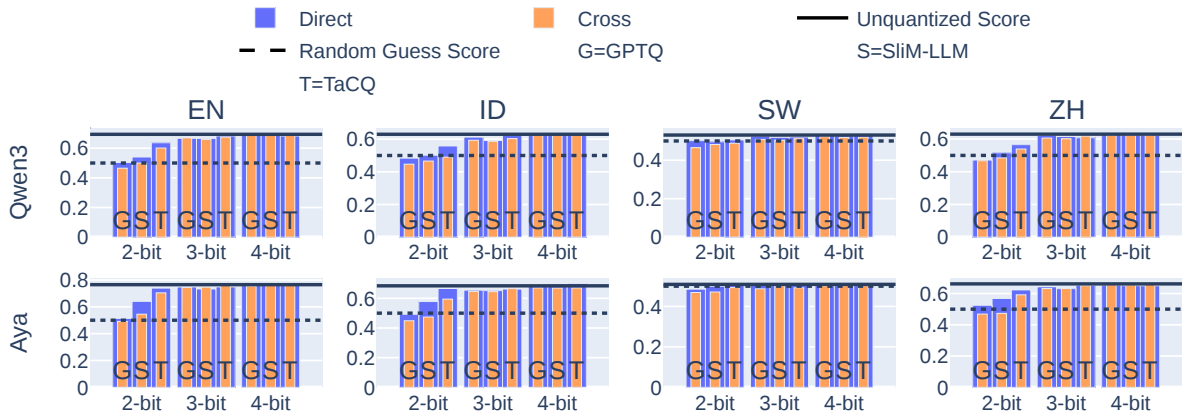

Figure 15: XStoryCloze Barplot

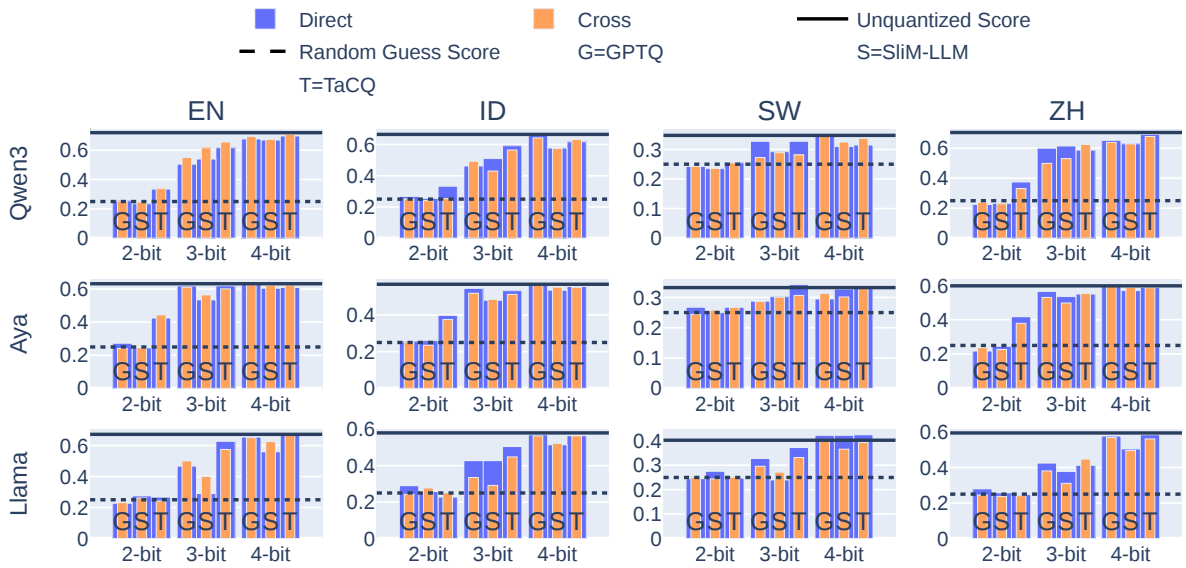

Figure 16: Global MMLU Lite Barplot

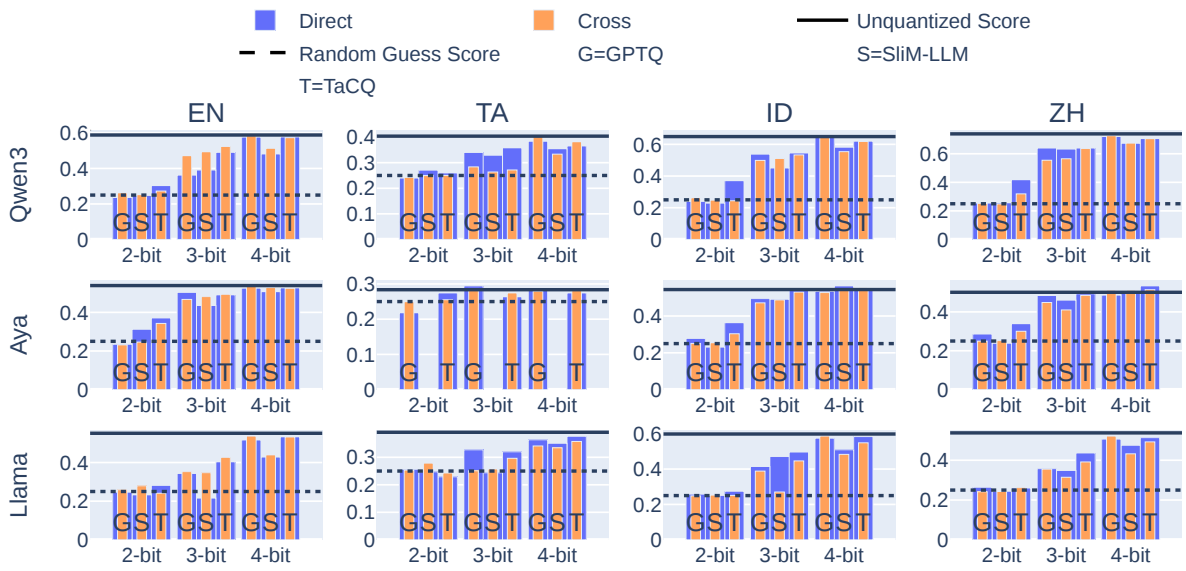

Figure 17: INCLUDE Barplot. INCLUDE does not have English evaluation, so we aggregate Dutch and German evaluations as a proxy for English.

Figure 18: Perplexity of Germanic (English) on cross-lingual calibration. The score is calculated using $C_{\text{target}}/\overline{C_{\text{cross}}}$ (lower is preferring English), aggregate across quantization methods.

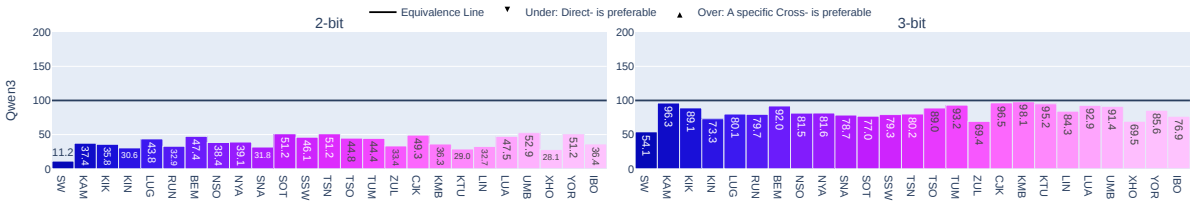

Figure 19: Perplexity of Niger-Congo (Swahili) on cross-lingual calibration. The score is calculated using $C_{\text{target}}/\overline{C_{\text{cross}}}$ (lower is preferring Swahili), aggregate across quantization methods.

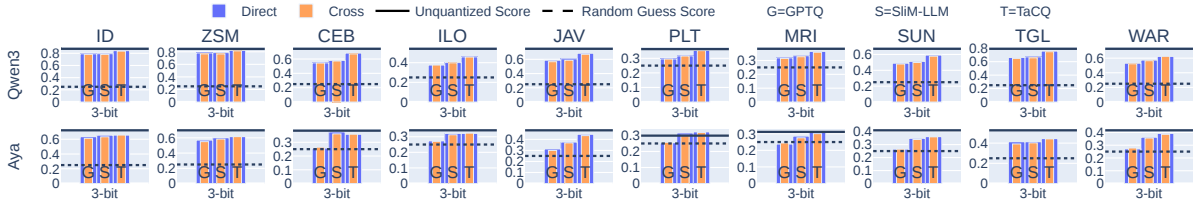

Figure 20: Belebele Austronesian (Indonesian) accuracy.

Quantization Technique - Calibration Language - Base Model

Eval Language		16-bit	Gptq					Slimllm					Tacq				
		-	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH
Qwen3	ID	0.87	-0.14	-0.13	-0.08	-0.10	-0.08	-0.14	-0.10	-0.08	-0.10	-0.12	-0.04	-0.07	-0.03	-0.06	-0.06
	ZSM	0.85	-0.13	-0.12	-0.06	-0.10	-0.10	-0.18	-0.12	-0.06	-0.11	-0.11	-0.04	-0.06	-0.02	-0.03	-0.05
	CEB	0.74	-0.25	-0.20	-0.20	-0.12	-0.24	-0.25	-0.17	-0.16	-0.14	-0.19	-0.10	-0.09	-0.06	-0.08	-0.11
	ILO	0.54	-0.18	-0.19	-0.17	-0.10	-0.19	-0.18	-0.13	-0.14	-0.13	-0.16	-0.09	-0.09	-0.08	-0.05	-0.13
	JAV	0.75	-0.28	-0.20	-0.16	-0.16	-0.26	-0.29	-0.18	-0.14	-0.16	-0.21	-0.13	-0.12	-0.06	-0.07	-0.11
	PLT	0.36	-0.09	-0.06	-0.07	-0.05	-0.08	-0.07	-0.02	-0.05	-0.05	-0.06	-0.01	-0.02	-0.01	-0.01	-0.02
	MRI	0.38	-0.11	-0.05	-0.07	-0.05	-0.10	-0.09	-0.05	-0.05	-0.08	-0.07	-0.05	-0.03	-0.02	-0.04	-0.03
	SUN	0.67	-0.26	-0.20	-0.18	-0.15	-0.25	-0.27	-0.18	-0.16	-0.15	-0.22	-0.17	-0.11	-0.08	-0.07	-0.14
	TGL	0.79	-0.21	-0.18	-0.13	-0.13	-0.18	-0.23	-0.12	-0.12	-0.13	-0.18	-0.08	-0.05	-0.03	-0.04	-0.09
	WAR	0.73	-0.26	-0.22	-0.20	-0.16	-0.25	-0.23	-0.18	-0.15	-0.15	-0.17	-0.09	-0.10	-0.10	-0.07	-0.14
Aya	ID	0.72	-0.11	-0.23	-0.09	-0.17	-0.12	-0.14	-	-0.07	-0.16	-0.16	-0.07	-0.04	-0.07	-0.09	-0.04
	ZSM	0.70	-0.15	-0.29	-0.13	-0.18	-0.15	-0.16	-	-0.10	-0.17	-0.17	-0.08	-0.06	-0.08	-0.11	-0.08
	CEB	0.38	-0.10	-0.12	-0.13	-0.08	-0.12	-0.06	-	-0.01	-0.06	-0.08	-0.03	-0.01	-0.03	-0.04	-0.01
	ILO	0.34	-0.09	-0.07	-0.07	-0.07	-0.09	-0.06	-	-0.02	-0.04	-0.06	-0.04	-0.01	-0.02	-0.00	-0.01
	JAV	0.48	-0.24	-0.24	-0.17	-0.16	-0.21	-0.13	-	-0.11	-0.11	-0.16	-0.09	-0.05	-0.04	-0.08	-0.08
	PLT	0.30	-0.04	-0.01	-0.04	-0.04	-0.05	-0.03	-	0.02	-0.05	-0.03	0.01	0.02	0.02	0.01	0.01
	MRI	0.31	-0.05	-0.04	-0.07	-0.06	-0.06	-0.07	-	-0.03	-0.06	-0.06	-0.03	-0.03	0.00	-0.05	-0.03
	SUN	0.41	-0.15	-0.16	-0.15	-0.14	-0.13	-0.09	-	-0.06	-0.10	-0.11	-0.07	-0.06	-0.05	-0.05	-0.04
	TGL	0.53	-0.15	-0.25	-0.11	-0.14	-0.18	-0.15	-	-0.11	-0.14	-0.18	-0.12	-0.09	-0.08	-0.08	-0.08
	WAR	0.42	-0.14	-0.13	-0.15	-0.09	-0.13	-0.09	-	-0.05	-0.07	-0.08	-0.06	-0.03	-0.02	-0.04	-0.06

Figure 21: Belebele Austronesian (Indonesian) accuracy: 3-bit, lighter or higher is better.

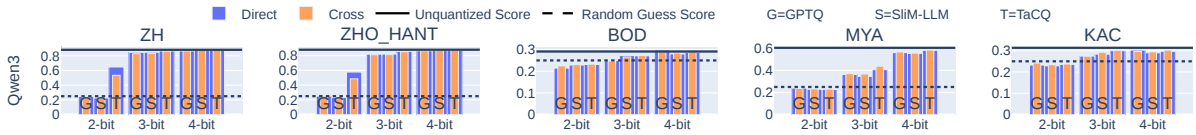

Figure 22: Belebele Sino-Tibetan (Chinese) accuracy. ZHO_HANT is Tradiational Chinese

Quantization Technique - Calibration Language - Base Model

Eval Language		16-bit	Gptq					Slimllm					Tacq				
		-	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH
Qwen3	ZH	0.88	-0.01	-0.01	-0.02	-0.01	-0.02	-0.01	-0.01	-0.01	-0.01	-0.01	-0.00	-0.01	-0.00	0.00	-0.00
	BOD	0.29	0.00	-0.00	-0.01	-0.01	0.01	-0.01	0.00	-0.01	0.00	-0.01	-0.01	0.00	0.00	0.01	-0.00
	KAC	0.31	-0.02	-0.03	-0.02	-0.02	-0.01	-0.02	-0.01	-0.02	-0.02	-0.02	0.00	-0.01	-0.01	-0.01	-0.02
	MYA	0.60	-0.03	-0.02	-0.06	-0.02	-0.04	-0.05	-0.03	-0.05	-0.04	-0.05	-0.04	-0.01	-0.01	-0.01	-0.02
	ZHO	0.88	-0.01	-0.03	-0.02	-0.02	-0.01	-0.02	-0.01	-0.01	0.00	-0.00	-0.01	-0.01	-0.00	-0.00	-0.01

Quantization Technique - Calibration Language - Base Model

Eval Language		16-bit	Gptq					Slimllm					Tacq				
		-	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH
Qwen3	ZH	0.88	-0.07	-0.07	-0.07	-0.08	-0.03	-0.05	-0.05	-0.06	-0.08	-0.03	-0.01	-0.04	-0.02	-0.02	-0.02
	BOD	0.29	-0.04	-0.05	-0.06	-0.04	-0.04	-0.04	-0.02	-0.04	-0.03	-0.02	-0.01	-0.02	-0.02	-0.02	-0.02
	KAC	0.31	-0.05	-0.04	-0.03	-0.04	-0.04	-0.03	0.01	-0.02	-0.01	-0.03	-0.03	-0.01	0.00	-0.01	-0.01
	MYA	0.60	-0.29	-0.22	-0.24	-0.16	-0.24	-0.28	-0.13	-0.22	-0.24	-0.26	-0.19	-0.09	-0.16	-0.11	-0.20
	ZHO	0.88	-0.06	-0.08	-0.08	-0.08	-0.06	-0.06	-0.08	-0.06	-0.09	-0.06	-0.04	-0.04	-0.02	-0.03	-0.02

Quantization Technique - Calibration Language - Base Model

Eval Language		16-bit	Gptq					Slimllm					Tacq				
		-	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH	EN	TA	ID	SW	ZH
Qwen3	ZH	0.88	-0.65	-0.64	-0.64	-0.64	-0.63	-0.65	-0.63	-0.66	-0.65	-0.66	-0.34	-0.51	-0.36	-0.63	-0.23
	BOD	0.29	-0.05	-0.08	-0.05	-0.05	-0.08	-0.04	-0.07	-0.06	-0.06	-0.06	-0.06	-0.04	-0.06	-0.06	-0.06
	KAC	0.31	-0.09	-0.05	-0.07	-0.04	-0.08	-0.08	-0.07	-0.07	-0.08	-0.08	-0.08	-0.05	-0.09	-0.08	-0.08
	MYA	0.60	-0.37	-0.36	-0.38	-0.38	-0.36	-0.37	-0.37	-0.36	-0.38	-0.38	-0.38	-0.36	-0.38	-0.38	-0.38
	ZHO	0.88	-0.62	-0.63	-0.64	-0.63	-0.63	-0.65	-0.64	-0.64	-0.65	-0.65	-0.37	-0.54	-0.40	-0.63	-0.30

Figure 23: Belebele Austronesian (Chinese) accuracy: 4-bit (Top), 3-bit (Middle), and 2-bit (Bottom), lighter or higher is better. ZHO is Traditional Chinese.