

# Quantization Meets Reasoning: Exploring LLM Low-Bit Quantization Degradation for Mathematical Reasoning

Anonymous ACL submission

## Abstract

Large language models have achieved significant advancements in complex mathematical reasoning benchmarks, such as MATH. However, their substantial computational requirements present challenges for practical deployment. Model quantization has emerged as an effective strategy to reduce memory usage and computational costs by employing lower precision and bit-width representations. In this study, we systematically evaluate the impact of quantization on mathematical reasoning tasks. Our results demonstrate that aggressive quantization methods like AWQ and GPTQ introduce up to 32.39% accuracy degradation (average 11.31%) on Llama-3 models, particularly in numerical computation and reasoning planning. To address this, we introduce a multidimensional evaluation framework combining qualitative capability analysis and quantitative error assessment. We further develop targeted recovery strategies, showing that fine-tuning quantized models on only 545 task-specific examples for 3 minutes on 4 GPUs effectively restores reasoning capabilities to near full-precision levels. Additionally, our error assessment pipeline achieves 98.9% accuracy in diagnosing and localizing errors across 3,366 failure cases, providing actionable insights for mitigating quantization-induced degradation.

## 1 Introduction

Large language models (LLMs) have substantially advanced the state of mathematical reasoning in artificial intelligence, demonstrating remarkable performance on diverse tasks ranging from basic arithmetic and quantitative reasoning to intricate geometric and competition-level problems (Brown et al., 2020; Chowdhery et al., 2023; Touvron et al., 2023; Achiam et al., 2023; OpenAI, 2023). Critically, these models excel not only at producing correct final answers, but also at providing step-by-step solutions that elucidate the underlying reasoning process (Lewkowycz et al., 2022; Wei et al.,

2022). Benchmarks such as MATH (Hendrycks et al., 2021) highlight these capabilities, where LLMs can guide humans through complex multi-step problems with detailed reasoning chains.

However, such advancements come at a cost. The computational requirements of LLMs, both in terms of memory and latency, pose significant practical barriers (Kaplan et al., 2020; Hoffmann et al., 2022; Gou et al., 2024). To address these efficiency challenges, researchers have explored model compression techniques such as pruning (LeCun et al., 1989; Han et al., 2015), knowledge distillation (Hinton, 2015; Jiao et al., 2019; Yang et al., 2024), and more recently, quantization (Hubara et al., 2018; Jacob et al., 2018; Yao et al., 2022). Quantization reduces memory usage and computational overhead by representing weights and activations in low-bit formats (e.g., INT8), halving GPU memory consumption and nearly doubling throughput in operations like matrix multiplication and attention (Rastegari et al., 2016; Lin et al., 2015; Dettmers et al., 2022). While it performs well on standard NLP tasks with minimal performance loss (Ma et al., 2024), its effect on complex mathematical reasoning, requiring precise, contextually coherent, and logical steps, remains unclear, particularly for tasks like MATH or Code.

Prior work has hinted at potential reasoning degradation under aggressive quantization (Shen et al., 2020; Kim et al., 2021; Lin et al., 2023), but a systematic understanding is lacking. Such limitations pose a stark contrast to the precision-based assumptions of advanced mathematical reasoning approaches. Models like Minerva (Lewkowycz et al., 2022) and reasoning strategies such as Chain-of-Thought (CoT) prompting (Wei et al., 2022; Xiong et al., 2023) rely heavily on high-fidelity internal representations to ensure logical consistency and correctness. Quantization, however, may disrupt these critical internal states. Meanwhile, ongoing developments in model optimization transcend

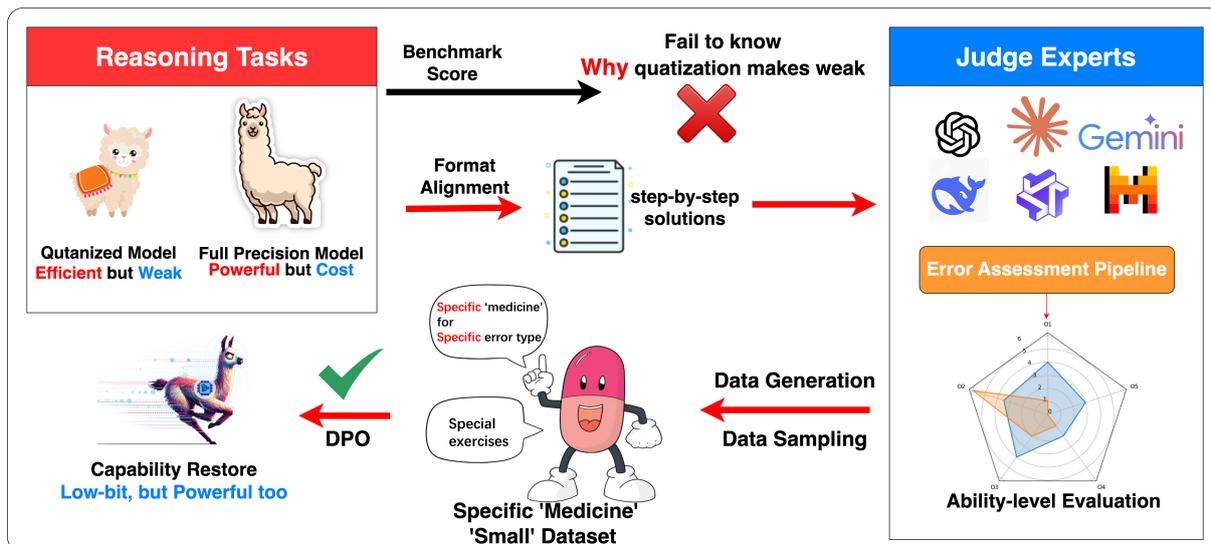


Figure 1: Pipeline of our study for evaluating and restoring reasoning capabilities in quantized models. By format alignment training and our error assessment pipeline with expert judge models, we diagnose and analyze ability-level reasoning weaknesses on model’s step-by-step solutions. Based on identified error types, we sample a targeted ‘medicine’ dataset to fine-tune the model via Direct Preference Optimization (DPO), aiming to restore performance while preserving efficiency.

simple training paradigms. Complex post-training pipelines that align models to human preferences and adapt them for specialized applications have gained traction, often involving intricate infrastructure and iterative refinement processes (Schulman et al., 2017; Rafailov et al., 2024).

Against this backdrop, the interplay between quantization, advanced inference strategies, and the underlying reasoning fidelity of LLMs emerges as a key research question. In this work, we seek to illuminate these relationships, providing insights into how quantization influences mathematical reasoning and how we might mitigate these effects.

To summarize, in this study, we present the following key contributions:

- We reveal that mainstream quantization methods (AWQ, GPTQ) incur substantial reasoning degradation, with up to 32.39% accuracy loss (average 11.31%) on MATH for Llama-3 models, exposing critical vulnerabilities in low-bit numerical representations.
- We propose a step-aligned evaluation protocol to dissect reasoning errors into four dimensions (conceptual, methodological, executional, logical), enabling granular diagnosis of quantization effects. Our automated error assessment pipeline achieves 98.9% accuracy in categorizing 3,366 failure cases, surpassing human-in-the-loop baselines.
- We demonstrate that quantized models recover

the performance gaps through lightweight fine-tuning on only 545 targeted examples, requiring only 3 minutes of training, thus enabling efficient deployment without sacrificing reasoning fidelity.

## 2 Related Work

### 2.1 Quantization Techniques

Modern quantization methods balance efficiency and performance through minimizing the model output differences after quantization. Post-training quantization (PTQ) approaches (Frantar et al., 2022; Lin et al., 2024; Xiao et al., 2023; Yao et al., 2022) enable efficient compression without retraining, while quantization-aware training (QAT) methods (Hu et al., 2021; Dettmers et al., 2024) preserve task-specific performance through learnable scaling factors. Recent work extends these techniques to LLMs, though primarily evaluated on language understanding rather than reasoning tasks.

### 2.2 Mathematical Reasoning in LLMs

The emergence of specialized models like Minerva (Lewkowycz et al., 2022) demonstrates LLMs’ potential for advanced mathematical problem-solving. Chain-of-Thought prompting (Wei et al., 2022) and its variants (e.g., Program-of-Thought (Chowdhery et al., 2023)) enhance multi-step reasoning by decomposing problems into interpretable sub-steps. However, these approaches assume high-precision model representations, po-

tentially conflicting with quantization’s reduced numerical precision.

## 2.3 Model Alignment and Reasoning Optimization

Recent advances in model alignment integrate instruction tuning, pioneered by the FLAN framework (Wei et al., 2021) and popularized through Alpaca (Taori et al., 2023), with preference optimization techniques like DPO (Rafailov et al., 2024). Concurrently, reasoning reliability enhancements employ self-consistency voting (Wang et al., 2022) and process-based reward models (Lightman et al., 2023a), building upon foundational work in verifiable reasoning (Creswell and Shanahan, 2022). Our investigation extends these directions by analyzing how quantization affects: the instruction-following capabilities crucial for step-by-step reasoning, and the consensus-building mechanisms in ensemble-based reasoning methods.

## 3 Methodology

### 3.1 Model Quantization

We conduct a comprehensive investigation into the effects of quantization techniques, examining both weight-only quantization methods (GPTQ (Frantar et al., 2022), AWQ (Lin et al., 2024)). Our evaluation encompasses various quantization configurations, specifically focusing on 4-bit weight precision with 16-bit activations (W4A16). Through the systematic application of these mainstream quantization techniques, we provide a rigorous and balanced analysis of the resulting quantized models, offering valuable insights into their performance characteristics and trade-offs.

### 3.2 Format Alignment

To address the challenge of inconsistent instruction following and irregular output formatting in model-generated solutions, we introduce a format alignment stage. This phase aims to instill in the model a structured, step-by-step reasoning workflow without altering its underlying mathematical knowledge. Crucially, the objective here is not to teach the model new mathematical facts, but rather to ensure strict adherence to a prescribed output format, thereby enabling reliable qualitative and quantitative analysis of reasoning capability across quantized and full-precision variants.

We employ LoRA (Hu et al., 2021) and QLoRA (Detmeters et al., 2024) for full-precision model and quantized model respectively as

## Training Prompt and Response Example

### System Prompt:

Your task is to solving mathematical problems by following these rules:

1. Follow the given instructions strictly and ensure the output aligns with the expected format.
2. Break down the reasoning process into clear, logical steps using numbered lists (e.g., 1., 2., 3.) to sequentially present each step. Each reasoning step should be isolated and clearly described to enhance readability and logical flow.
3. After completing all reasoning steps, provide the final answer.
4. Please reason step by step, and put your final answer within `\boxed{}`. The final answer must be formatted as `\boxed{your_final_answer}`. For example: Answer: `\boxed{42}`.
5. Maintain consistent formatting across training and evaluation phases to ensure the model outputs meet evaluation requirements seamlessly.

### Problem Sample (from PRM-800K Dataset Training Subsets):

The product of the proper positive integer factors of  $n$  can be written as  $n^{\frac{ax+b}{c}}$ , where  $x$  is the number of positive divisors  $n$  has,  $c$  is a positive integer, and the greatest common factor of the three integers  $a$ ,  $b$ , and  $c$  is 1. What is  $a+b+c$ ? [Gold Answer: 1]

### Response Sample (from PRM-800K Dataset Training Subsets):

1. I notice that the problem involves the product of the proper positive integer factors of  $n$ , which means all the factors of  $n$  except  $n$  itself.
2. I also notice that the problem involves the number of positive divisors of  $n$ , which I'll call  $x$ .
3. I wonder if there is a way to relate the product of the proper factors of  $n$  to  $n$  and  $x$ .
4. I recall that the product of all the factors of  $n$  is  $n^{\frac{x+1}{2}}$ , since each factor can be paired with its complement, and their product is  $n$ .
5. Therefore, the product of the proper factors of  $n$  is  $n^{\frac{x}{2}}$ , which is  $n^{\frac{(x-2)}{2}}$ .
6. This means that the product of the proper factors of  $n$  can be written as  $n^{\frac{(ax+b)}{c}}$ , where  $a=1$ ,  $b=-2$ , and  $c=2$ .
7. The greatest common factor of  $a$ ,  $b$ , and  $c$  is 1, as required.
8. The answer is  $a+b+c=1+(-2)+2=1$ . Answer: `\boxed{1}` [Answer: 1]

Figure 2: An example of the data used during training. The output is concatenated using added tokens, and the final answer is filled in `\boxed{}` in the evaluation format.

lightweight adaptation techniques for format alignment. By selectively updating a small subset of parameters, these methods efficiently align knowledge of step-by-step solution formats into the model’s latent space without extensive retraining. This fine-tuning enables us to observe how multi-step reasoning is preserved or altered once the model is quantized, offering deeper insights into any capability loss induced by compression.

For training, we utilize the PRM800K dataset (Lightman et al., 2023b), which provides 800K step-level correctness annotations from 75K solutions to 12K problems. These annotations supply granular, step-by-step reasoning trajectories, equipping models to separate complex problem-solving processes into well-defined stages. To reinforce this structure, we adopt a consistent system prompt across training and evaluation, ensuring that the boundaries of logical steps and final answers are clearly delineated. This consistent, step-by-step alignment is a necessary foundation for our subsequent qualitative and

MATH		Vanilla (BF16)	AWQ (W4A16)	GPTQ (W4A16)
Llama-3.1-8B-Instruct	score	47.2	41.8	43.06
	$\Delta$	-	$\downarrow 5.4(11.44\%)$	$\downarrow 5.8(12.29\%)$
Llama-3.2-3B-Instruct	score	40.2	38.8	31.6
	$\Delta$	-	$\downarrow 2.16(4.97\%)$	$\downarrow 4.0(10.42\%)$
Llama-3.2-1B-Instruct	score	18.6	14.26	14.94
	$\Delta$	-	$\downarrow 2.8(16.67\%)$	$\downarrow 6.8(32.39\%)$

Table 1: Evaluation results of full-precision models and 4-bit weight, 16-bit activation quantization methods (AWQ and GPTQ) on the MATH benchmark. All three model scales exhibit varying degrees of performance degradation, posing challenges in solving mathematical reasoning problems. Notably, smaller-scale models suffer a performance drop exceeding 30%.

quantitative analyses of quantization-induced degradation in mathematical reasoning.

### 3.3 Analyses for Reasoning Errors

To investigate the underlying reasons for the performance degradation observed in the quantized models, we categorized the errors into seven specific error types in four major error categories.:

- **Conceptual Errors**

- **Conceptual misunderstanding:** Inadequate grasp of fundamental concepts or principles, leading to an incorrect framing of the problem or misguided approach to the solution.
- **Contextual oversight:** Failure to account for relevant contextual factors or domain constraints (e.g., physical limits, geometric considerations) that significantly affect the problem’s outcome.

- **Method Errors**

- **Procedural error:** Errors in executing or adhering to prescribed procedures (e.g., failure to follow a standard algorithm or overlooking a required step), resulting in incomplete or invalid solutions.
- **Formula rule error:** Improper application of mathematical rules, theorems, or formulae (including misapplication of a formula to an inappropriate scenario), which undermines the validity of the derived results.

- **Execution Errors**

- **Computational error:** Mistakes in arithmetic or algebraic manipulation, such as incorrect summation, factorization, or expansion, that compromise the accuracy of the final answer.
- **Symbolic manipulation error:** Inaccurate handling of symbols or expressions, includ-

ing mislabeling variables or misinterpreting symbolic transformations, leading to an incorrect representation of the problem.

- **Reasoning Errors**

- **Logical reasoning error:** Breakdown in the logical chain of thought, where the inference steps do not follow coherently or omit essential logical links, causing a disconnect between premises and conclusion.

### 3.4 Capability Restore

Building upon our systematic error analysis, we develop targeted capability restoration strategies to address quantization-induced performance degradation. For each identified error category (conceptual, methodological, executional, and logical), we create focused training subsets by pairing failed step-wise responses with correct responses to form a dataset of positive and negative pairs:

- **Conceptual Restoration:** High-context problems requiring domain-specific constraints.
- **Methodological Restoration:** Multi-step derivations with annotated procedure adherence.
- **Executional Restoration:** Computation-intensive problems with stepwise verification.
- **Logical Restoration:** Chain-of-thought reasoning with explicit inference links.

We implement Direct Preference Optimization (DPO) with parameter-efficient adaptation, using QLoRA for quantized models. This dual approach enables simultaneous optimization of two objectives: (1) maximizing the likelihood of correct solution trajectories through supervised fine-tuning on error-specific subsets, and (2) addressing the model’s capability gaps while leveraging the undiminished foundational abilities of the quantized model to complete the solution, thereby enhancing overall mathematical reasoning performance.

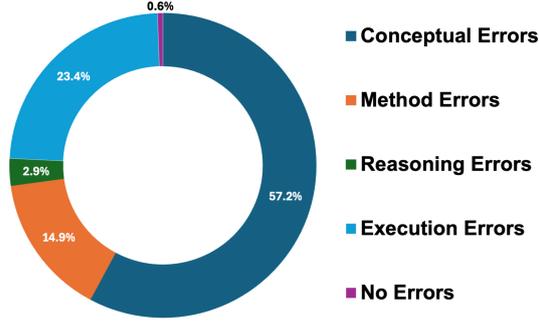


Figure 3: Distribution of error types in quantized models, highlights the dominant error types affecting mathematical reasoning in quantized models.

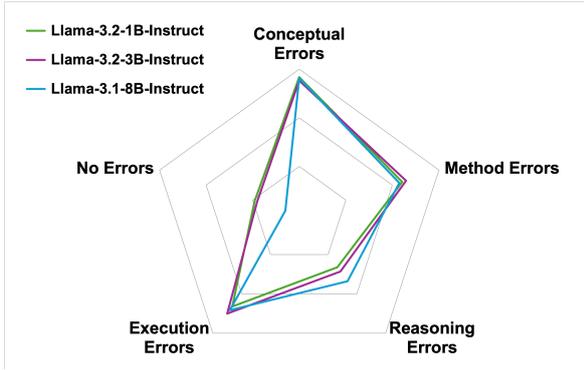


Figure 4: The radar plot of error distributions across different scale Llama models on the MATH benchmark.

Our recovery process comprises three phases:

1. **Data Extraction and Construction:** We extract and construct a dataset from the quantized models’ failed responses across different model scales.
2. **Data Deduplication and Cross-Scale Sampling:** To avoid redundancy, we deduplicate the dataset and perform cross-scale sampling based on error rates. This facilitates cross-error joint training with a mix of problem difficulties.
3. **DPO Training on the Aligned Model:** Building on the format-aligned model, we further train using a DPO approach.

This phased approach achieves significant capability recovery while introducing less than 1% additional parameter updates relative to the original model sizes. Our experiments demonstrate that targeted restoration using as few as 545 error-specific examples per category can bridge the performance gap between quantized and full-precision models. Remarkably, even training with a dataset comprising only 323 samples for a single error type can restore up to 9% of the model’s capability, confirming that localized capability losses can be efficiently remediated through structured retraining.

---

### Algorithm 1 Pseudocode of our approach.

---

**Input:** Full-precision model  $M$ , Quantization method  $Q$ , Step-annotated dataset  $D_{train}$ , Evaluation dataset  $D_{test}$ , Expert models  $\{E_1 \dots E_n\}$

**Output:** Enhanced quantized model  $M_{enhanced}$

- 1: **Step 1: Model Quantization**
  - 2: Apply quantization method  $Q$  to model  $M$ :
  - 3:  $M_Q \leftarrow Q(M)$
  - 4: **Step 2: Format Alignment**
  - 5: Fine-tune  $M_Q$  with QLoRA on  $D_{train}$ :
  - 6:  $M_A \leftarrow \text{QLoRA}(M_Q, D_{train})$
  - 7: **Step 3: Error Evaluation**
  - 8: Generate outputs  $O$  on  $D_{test}$ :
  - 9:  $O \leftarrow M_A(D_{test})$
  - 10: Initialize error set  $E \leftarrow \emptyset$
  - 11: **for** each output  $o_i \in O$  **do**
  - 12:   **Multi-Expert Evaluation:**
  - 13:   Collect evaluations  $\{E_1(o_i) \dots E_n(o_i)\}$
  - 14:   **if** Majority vote  $\neq$  DeepSeek-R1 **and** votes  $< 4$  **then**
  - 15:     Flag  $o_i$  as Rejected
  - 16:      $E \leftarrow E \cup \{\text{ErrorType}(o_i)\}$
  - 17:   **end if**
  - 18: **end for**
  - 19: **Step 4: Human Verification**
  - 20: Random sample  $s \leftarrow 2\%$  of  $E$ :
  - 21:  $E_{sample} \leftarrow \text{RandomSample}(E, s)$
  - 22:  $E_{verified} \leftarrow \text{HumanVerify}(E_{sample})$
  - 23: **Step 5: Capability Restoration**
  - 24: **for** each error type  $t \in E_{verified}$  **do**
  - 25:   Generate targeted dataset  $D_t$ :
  - 26:    $D_t \leftarrow \text{CurateData}(t, D_{train})$
  - 27:   Apply DPO training:
  - 28:    $M_A \leftarrow \text{DPO}(M_A, D_t)$
  - 29: **end for**
  - 30: **Return:**  $M_{enhanced} \leftarrow M_A$
- 

## 4 Experiments 314

### 4.1 General Setup 315

All experiments were conducted on a cluster comprising 4 Nvidia 80G A800 GPUs. 316

#### Base Model & Scale 317

Given the general applicability of our approach to mainstream models and the versatility across different model sizes, we select three distinct models from the Llama 3 series as the foundation for our experiments. Specifically, we choose models of varying sizes, namely 8B for Llama-3.1 and 1B/3B for Llama-3.2, all Instruct version. 318

For all of these models, we observe significant 319

math accuracy degradation resulting from quantization, and we also demonstrate the powerful effect of our proposed fast reconstruction method.

**Quantization Strategy.** Our experiments employ two mainstream quantization techniques: AWQ and GPTQ (Frantar et al., 2022; Lin et al., 2024), both known for their strong performance in model compression. These methods have led to numerous low-precision models widely adopted in practice. Our work aims to provide complementary insights for community.

**Training Strategy.** We use a sequence length of 1024 and in both the format alignment and capability restoration stages, we employed a global batch size of 16 and a warm-up period spanning 3%. All models were trained for 3 epochs.

**Evaluation Strategy.** We evaluated our models using the OpenCompass evaluation framework (Contributors, 2023), maintaining consistency in both the global batch size and the prompt with those used during training. This setup ensures a fair comparison across all models.

## 4.2 Format Alignment

**Setting.** We applied LoRA fine-tuning to the vanilla model, while QLoRA (Dettmers et al., 2024) was employed for the quantized models. The LoRA rank was set to 32 for the format alignment stage. During training, the system prompt strictly enforced a step-wise output of the reasoning process, and outputs that deviated from the required format were penalized.

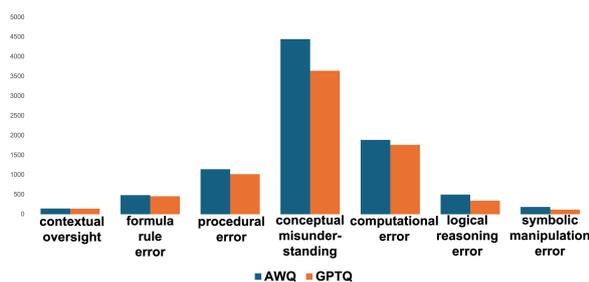


Figure 5: Comparative analysis of error types for quantized models (AWQ-W4A16 and GPTQ-W4A16) across three model scales.

## 4.3 Analyses for Reasoning Errors

**Error Assessment Pipeline.** To facilitate a rigorous and scalable evaluation of quantization-induced errors in reasoning tasks, we developed an automated assessment pipeline powered by state-of-the-art language models. This pipeline systematically processes model outputs and classifies errors according to our predefined *error\_types\_list* taxon-

omy. By leveraging a pre-trained transformer as the core evaluator, we reduce subjective bias and ensure consistent, reproducible error analyses across all experimental conditions. Furthermore, the computational scoring framework supports high-throughput performance assessment while preserving granularity in error categorization.

Our evaluation pipeline comprises three primary stages:

- Model Judge:** For each instance in which a quantized model produces an incorrect answer, we employ a dedicated “judge model” to analyze the error. The judge model is tasked with: (a) identifying the first occurrence of an error, (b) specifying the exact step where the error is introduced, (c) assigning an error category based on a nested classification scheme, and (d) providing an explanation along with a confidence score for its determination.
- Majority Vote:** To mitigate hallucinations and enhance evaluation stability, we implement a majority voting mechanism across five models: DeepSeek-R1, GPT-4o, GPT-4, Qwen-Max, and DeepSeek-V3. DeepSeek-R1’s assessment is treated as the baseline, and its results are compared with the majority vote from the other models. Instances of disagreement are flagged for further review, ensuring consistency and minimizing spurious judgments.
- Human Verification:** For cases with conflicting assessments from the majority vote, manual review is conducted. Additionally, we randomly sample 2% of the evaluated cases to verify the accuracy and consistency of the automated judgments.

This pipeline generally aligns with human error analysis and minimizes misclassifications and inconsistencies. However, we also encountered unexpected scenarios. For instance, when the canonical answer is  $\frac{1}{2}$  but the quantized model outputs “5.5” the judge model sometimes erroneously concludes that there is “No Error” due to subtle discrepancies in reasoning or formatting even if we do not give this type of judgment. Such findings underscore both the robustness of our judge framework.

**Quantitative analyses.** Table 1 presents a performance comparison between full-precision and

MATH-500			Vanilla (BF16)	AWQ (W4A16)	GPTQ (W4A16)
Llama-3.1-8B-Instruct	Aligned	score	37.36	30.4	30.8
		$\Delta$	-	$\downarrow 6.94(18.58\%)$	$\downarrow 6.56(17.56\%)$
Llama-3.2-3B-Instruct	Aligned	score	37.4	34.4	31.4
		$\Delta$	-	$\downarrow 3.0(8.02\%)$	$\downarrow 6.0(16.04\%)$
Llama-3.2-1B-Instruct	Aligned	score	21	14.2	14.94
		$\Delta$	-	$\downarrow 6.8(32.39\%)$	$\downarrow 6.2(29.52\%)$

Table 2: Evaluation results of format-aligned full-precision models and quantization methods (AWQ and GPTQ) on the MATH-500 benchmark. After alignment, the model is required to output responses strictly following the numbered step-wise format, resulting in slight changes in evaluation score.

quantized models on the MATH dataset across three model scales. We report both the performance (Score) and the corresponding accuracy degradation ( $\Delta$ ) induced by quantization. The vanilla model serves as the upper baseline, while quantized models exhibit varying degrees of performance loss. Notably, our results indicate that the 1B model is particularly sensitive to quantization.

Based on the set of questions answered incorrectly by the quantized models, we conducted further quantitative analyses. Figure 3 and 5 illustrates the overall distribution of the four error categories. Conceptual errors account for the majority, representing more than half of all observed mistakes. Execution errors follow next. Our case study analyses reveal that quantized models are particularly deficient in basic arithmetic operations, with multiplication errors showing the most significant performance decline. We hypothesize that these computation errors, especially in models quantized to 4-bit weights, are likely attributable to low-bit precision issues such as overflow and underflow, which propagate inaccuracies throughout multi-step calculations. Additional illustrative cases are provided in the Appendix.

Intuitively, one might expect that reasoning errors would be more frequent given the inherent complexity of mathematical reasoning. However, our analysis shows that reasoning errors are less common. This observation may reflect a “survivor effect,” whereby earlier errors, such as logical or computational failures, prevent the model from reaching subsequent steps where boundary conditions would be tested. For precise statistics, we recorded only the first error type encountered in each instance. These findings suggest that quantization disrupts the step-by-step reasoning process, thereby exposing the models’ limitations in managing complex mathematical tasks. The underlying

causes likely stem from limited numerical precision and the accumulation of quantization noise in low-bit formats, which impair the representation of long-term dependencies and intermediate values during complex reasoning.

We hypothesize that small losses in specific capability dimensions can be compensated by fine-tuning on a minimal subset of targeted data. This approach could potentially restore overall performance by addressing the specific weaknesses introduced by quantization. For instance, targeted training on multi-step reasoning tasks could reduce step omissions by reinforcing the model’s ability to retain intermediate results. By strategically addressing these specific failure modes, it may be possible to recover and enhance the overall capabilities of quantized models, leveraging their efficiency while minimizing performance degradation.

#### 4.4 Capability Restore

**Dataset Sampling.** We constructed the dataset for capability restoration by sampling from all MATH problems that were answered correctly by the full-precision model but incorrectly by the quantized model. Data sampling was guided by proportional representation and consistency with expert model major vote outcomes. To enhance data validity, we performed deduplication and prioritized cases where the major vote results were most consistent and the errors were more pronounced. We aimed to maintain a compact dataset to reduce training resource consumption. Ultimately, we effectively sampled 545 examples from 3329 quantized model failure cases. The full-precision model’s correct answers were used as positive examples, while the quantized model’s incorrect answers served as negative examples, forming our DPO dataset.

**Setting.** Training for capability restoration stage was conducted on the format-aligned model. For

Full Precision	Llama-3.2-1B-Instruct		Llama-3.2-3B-Instruct	
	18.6		40.2	
	AWQ	GPTQ	AWQ	GPTQ
<b>Aligned</b>	<b>14.26</b>	<b>14.94</b>	<b>38.8</b>	<b>31.6</b>
<b>Setting 0</b>	17.2	18.6	41.6	38.2
<b>Setting 1</b>	15.8	18.0	<b>41.2</b>	36.2
<b>Setting 2</b>	<b>16.6</b>	16.6	37.2	34.4
<b>Setting 3</b>	15.2	15.4	35.2	33.0

Table 3: Ablation study settings and evaluation results on the MATH-500 dataset. Setting 0: All error cases; Setting 1: Conceptual error cases only; Setting 2: Methodological error cases only; Setting 3: Execution error cases only.

this stage, the LoRA rank was still set to 32, and a cosine learning rate schedule was employed with a warm-up ratio of 0.1. We used a sigmoid function as the preference loss function, with a global batch size of 8 and a learning rate of  $1 \times 10^{-6}$ . For evaluation, to ensure fairness and mitigate data leakage risks, we used the MATH-500 dataset. MATH-500 comprises 500 problems that are not included in the PRM800K dataset and were randomly sampled from the MATH dataset, thereby ensuring a similar difficulty distribution.

Full Precision	Llama-3.2-1B-Instruct		Llama-3.2-3B-Instruct	
	21		37.4	
	AWQ	GPTQ	AWQ	GPTQ
<b>Format Aligned</b>	14.26	14.94	38.8	31.6
<b>Restored</b>	18.2	18.8	41.6	38.2

Table 4: Capability restoration experiment results on the MATH-500 benchmark. Using only 545 samples within a few minutes.

**Results.** We conducted rapid capability restoration training on two smaller-scale models. Table 4 presents the restoration results using a dataset of 545 sampled examples. Due to the limited data volume, training required only 2–4 minutes on 4 GPUs, demonstrating minimal resource consumption. Notably, the quantized models exhibited marked improvements on this dataset.

#### 4.5 Ablation Study

We conducted quantitative analyses focusing on the earliest error type observed during each model’s reasoning process. We hypothesize that quantized models may exhibit a “bucket effect” in reasoning tasks, wherein they still perform well on instruction following yet fail to solve complex reasoning problems. To investigate this, we isolated the three most frequently occurring early error types for ablation experiments: (1) Conceptual error cases only, (2) Method error cases only, (3) Execution errors only. (4) All error types cases. This targeted analysis allowed us to assess the specific contributions of

each error category to the overall degradation in reasoning performance, thereby providing insights into how quantization selectively impacts different aspects of model capability.

The ablation experiments results reveal that even when training on datasets containing only a single type of error, the model’s capabilities can be significantly improved. For instance, using only data associated with conceptual errors enabled the quantized model to recover reasoning performance comparable to that of the full-precision model. Notably, this subset comprised merely 323 examples and takes very few minutes for training, which required minimal computational resources. These findings lend support to our hypothesis: quantization does not diminish the model’s fundamental abilities but selectively degrades specific dimensions that are crucial for solving reasoning tasks, such as those encountered in mathematical problem solving.

## 5 Conclusion

This study investigates the performance degradation of low-bit quantized LLMs in complex mathematical reasoning tasks. We find that quantization disproportionately impacts conceptual and computational reasoning, primarily due to precision loss and noise propagation in low-bit representations. To address these issues, we introduce a multidimensional evaluation pipeline and demonstrate that targeted capability restoration strategies can significantly recover performance with minimal training using our specific ‘medicine’ datasets. Our parameter-efficient adaptation methods, such as QLoRA and DPO, offer practical solutions for resource-constrained deployments. These findings highlight the trade-off between efficiency and reasoning fidelity, providing actionable insights for optimizing quantized LLMs. Future work will explore extending this framework to other reasoning domains and investigating potential advantages of reduced precision in specific scenarios.

## 6 Limitations & Future Works

In this study, we investigated the performance degradation of quantized models on the competitive mathematical reasoning benchmark MATH through both qualitative and quantitative analyses, employing seven distinct capability dimensions for evaluation. However, our work is constrained by limited time and computational resources, which restrict the scale of our experiments and prevent comprehensive case studies or ablation analyses across models of varying sizes. Additionally, we did not explore whether the planning failures and computational errors observed in low-bit models are caused by storage overflows resulting from low-precision data formats through knowledge exposure experiments.

Moreover, we identified intriguing phenomena where quantized models outperform full-precision models in certain specific mathematical solving steps. However, we did not conduct extensive experiments to further investigate these enhancements. As a short-term objective, we plan to perform more targeted analyses on the degraded capabilities of quantized models and explore efficient methods for their recovery, incorporating more human involvement in case studies to uncover deeper insights. Furthermore, we aim to extend our evaluation to other types of reasoning tasks, assessing quantized models from the perspective of their inherent properties to achieve a more comprehensive understanding of their performance across diverse reasoning scenarios.

## References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.

OpenCompass Contributors. 2023. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>.

Antonia Creswell and Murray Shanahan. 2022. Faithful reasoning using large language models. *arXiv preprint arXiv:2208.14271*.

Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35:30318–30332.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2024. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36.

Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*.

Zhibin Gou, Zhihong Shao, Yeyun Gong, yelong shen, Yujie Yang, Minlie Huang, Nan Duan, and Weizhu Chen. 2024. ToRA: A tool-integrated reasoning agent for mathematical problem solving. In *The Twelfth International Conference on Learning Representations*.

Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.

Geoffrey Hinton. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. 2022. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.

Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2018. Quantized neural networks: Training neural networks with low precision weights and activations. *Journal of Machine Learning Research*, 18(187):1–30.

667	Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. 2018. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In <i>Proceedings of the IEEE conference on computer vision and pattern recognition</i> , pages 2704–2713.	721
668		722
669		723
670		724
671		725
672		
673		
674	Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. Tinybert: Distilling bert for natural language understanding. <i>arXiv preprint arXiv:1909.10351</i> .	727
675		728
676		729
677		730
678	Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. <i>arXiv preprint arXiv:2001.08361</i> .	731
679		732
680		733
681		734
682		735
683	Sehoon Kim, Amir Gholami, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. 2021. I-bert: Integer-only bert quantization. In <i>International conference on machine learning</i> , pages 5506–5518. PMLR.	736
684		
685		
686		
687	Yann LeCun, John Denker, and Sara Solla. 1989. Optimal brain damage. <i>Advances in neural information processing systems</i> , 2.	737
688		738
689		739
690	Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. 2022. Solving quantitative reasoning problems with language models. <i>Advances in Neural Information Processing Systems</i> , 35:3843–3857.	740
691		741
692		742
693		743
694		744
695		745
696		746
697	Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023a. <a href="#">Let’s verify step by step</a> . <i>Preprint</i> , arXiv:2305.20050.	747
698		748
699		749
700		750
701		751
702	Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023b. <a href="#">Let’s verify step by step</a> . <i>arXiv preprint arXiv:2305.20050</i> .	752
703		753
704		754
705		755
706		756
707	Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Weiming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. <i>Proceedings of Machine Learning and Systems</i> , 6:87–100.	757
708		758
709		759
710		760
711		761
712		762
713	Jiaju Lin, Haoran Zhao, Aochi Zhang, Yiting Wu, Huqiyue Ping, and Qin Chen. 2023. <a href="#">Agentsims: An open-source sandbox for large language model evaluation</a> . <i>ArXiv</i> , abs/2308.04026.	763
714		764
715		765
716		766
717	Zhouhan Lin, Matthieu Courbariaux, Roland Memisevic, and Yoshua Bengio. 2015. Neural networks with few multiplications. <i>arXiv preprint arXiv:1510.03009</i> .	767
718		768
719		769
720		770
		771
		772
		773
		774
		775
		776
	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. <i>arXiv preprint arXiv:2402.17764</i> .	
	OpenAI. 2023. <a href="#">Gpt-4v(ision) system card</a> .	
	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. <i>Advances in Neural Information Processing Systems</i> , 36.	
	Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In <i>European conference on computer vision</i> , pages 525–542. Springer.	
	John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. <i>arXiv preprint arXiv:1707.06347</i> .	
	Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. 2020. Q-bert: Hessian based ultra low precision quantization of bert. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 34, pages 8815–8821.	
	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Stanford alpaca: An instruction-following llama model.	
	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. <i>arXiv preprint arXiv:2307.09288</i> .	
	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. <i>arXiv preprint arXiv:2203.11171</i> .	
	Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. <i>arXiv preprint arXiv:2109.01652</i> .	
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	
	Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In <i>International Conference on Machine Learning</i> , pages 38087–38099. PMLR.	

777       Jing Xiong, Zixuan Li, Chuanyang Zheng, Zhijiang  
778       Guo, Yichun Yin, Enze Xie, Zhicheng Yang, Qingx-  
779       ing Cao, Haiming Wang, Xiongwei Han, et al. 2023.  
780       Dq-lore: Dual queries with low rank approximation  
781       re-ranking for in-context learning. *arXiv preprint*  
782       *arXiv:2310.02954*.

783       Runming Yang, Taiqiang Wu, Jiahao Wang, Pengfei  
784       Hu, Ngai Wong, and Yujiu Yang. 2024. Llm-neo:  
785       Parameter efficient knowledge distillation for large  
786       language models. *arXiv preprint arXiv:2411.06839*.

787       Zhewei Yao, Reza Yazdani Aminabadi, Minjia Zhang,  
788       Xiaoxia Wu, Conglong Li, and Yuxiong He. 2022.  
789       Zeroquant: Efficient and affordable post-training  
790       quantization for large-scale transformers. *Advances*  
791       *in Neural Information Processing Systems*, 35:27168–  
792       27183.

793  
794  
795  
796  
797  
798  
799  
800  
801  
802  
803  
804  
805  
806  
807

## A Solutions by Quantization Models

### A.1 Execution Errors: Computational error

In Case #93, the task was to find the value of  $c$  for which the circle defined by  $x^2 - 10x + y^2 + 6y + c = 0$  has a radius of 1. The correct approach involves completing the square, resulting in  $(x - 5)^2 + (y + 3)^2 = -c + 34$ , and solving  $-c + 34 = 1$  to find  $c = 33$ . However, the GPTQ-Model made a computational error by incorrectly rearranging the equation as  $34 - c = 1$  leading to the wrong answer  $c = 35$ . This error stemmed from mishandling the sign of  $c$  during algebraic manipulation, highlighting the importance of careful sign management in equation solving. The case #93 shows the **Computational Error**.

#### Problem #93:

For what value of  $c$  will the circle with equation  $x^2 - 10x + y^2 + 6y + c = 0$  have a radius of length 1? [Gold Answer: 33]

#### Response of LLaMA-3.1-8B-Instruct-BF16 (Vanilla):

- To find the radius of a circle, I need to complete the square for both  $x$  and  $y$  terms in the equation. \n
- I can rewrite the equation as  $(x^2 - 10x) + (y^2 + 6y) = -c$ . \n
- To complete the square for  $x^2 - 10x$ , I need to add  $(10/2)^2 = 25$  to both sides. \n
- To complete the square for  $y^2 + 6y$ , I need to add  $(6/2)^2 = 9$  to both sides. \n
- This gives me  $(x^2 - 10x + 25) + (y^2 + 6y + 9) = -c + 25 + 9$ . \n
- I can factor the left side as  $(x - 5)^2 + (y + 3)^2$ . \n
- The right side simplifies to  $-c + 34$ . \n
- So the equation becomes  $(x - 5)^2 + (y + 3)^2 = -c + 34$ . \n
- Now I can see that the radius of the circle is the square root of the right side. \n
- I want the radius to be 1, so I need to have  $-c + 34 = 1^2 = 1$ . \n
- Solving for  $c$ , I get  $c = 34 - 1 = 33$ . [Answer: 33]

#### Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQModel):

- To find the radius of a circle, I need to complete the square for both the  $x$  and  $y$  terms in the equation and write it in the standard form  $(x - h)^2 + (y - k)^2 = r^2$ , where  $(h, k)$  is the center of the circle and  $r$  is the radius. \n
- I will start with the  $x$  terms and try to make a perfect square by adding and subtracting the same constant. \n
- I notice that the coefficient of  $x$  is  $-10$ , so I want to add and subtract  $(-10/2)^2 = 25$  to get  $x^2 - 10x + 25 - 25$ . \n
- I can factor the perfect square as  $(x - 5)^2$  and simplify the equation to  $(x - 5)^2 - 25 + y^2 + 6y + c = 0$ . \n
- I will do the same thing for the  $y$  terms and add and subtract  $(-6/2)^2 = 9$  to get  $y^2 + 6y + 9 - 9$ . \n
- I can factor the perfect square as  $(y + 3)^2$  and simplify the equation to  $(x - 5)^2 - 25 + (y + 3)^2 - 9 + c = 0$ . \n
- Now I have the equation in the form  $(x - 5)^2 + (y + 3)^2 - 34 + c = 0$ . \n
- To make the radius 1, I need to have  $r^2 = 1$ , so I will set  $-34 + c = 1$  and solve for  $c$ . \n
- Adding 34 to both sides, I get  $c = 35$ . [Answer: 35]

Figure 6: Problem #93: Computational Error.

### A.2 Reasoning Errors: Logical reasoning error

In Problem #128, the GPTQModel made a logical reasoning error when calculating the minimum number of miles Suzanne could walk in February. While the model correctly identified that February has 28 or 29 days, it missed the 27th day when calculating her walking schedule every third day, leading to an incorrect conclusion of 8 walking days and 32 miles. In reality, the correct number of walking days is 9 (3, 6, 9, 12, 15, 18, 21, 24, 27), resulting in a total of  $9 \times 4 = 36$  miles. This error highlights the model's logical reasoning gap in iterating through sequential intervals accurately. The case #128 shows the **Logical Reasoning Error**.

808  
809  
810  
811  
812  
813  
814  
815  
816  
817  
818  
819  
820  
821

#### Problem #128:

Suzanne walks four miles every third day. What is the fewest number of miles she can walk in February? [Gold Answer: 36]

#### Response of LLaMA-3.1-8B-Instruct-BF16 (Vanilla):

- To find the fewest number of miles Suzanne can walk in February, I need to figure out how many days she walks four miles in that month. \n
- February has 28 or 29 days, depending on whether it is a leap year or not. \n
- I don't know if this is a leap year or not, so I will assume it has 28 days. \n
- If Suzanne walks four miles every third day, that means she walks four miles on the 3rd, 6th, 9th, 12th, 15th, 18th, 21st, 24th, and 27th of February. \n
- That's nine days, so she walks a total of  $9 \times 4 = 36$  miles in February. \n
- This is the fewest number of miles she can walk, because if she didn't walk on any of those days, she would have to walk on another day, which would increase her total miles. \n
- Therefore, the answer is 36 miles. [Answer: 36]

#### Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQModel):

- To find the fewest number of miles Suzanne can walk in February, I need to figure out how many days she walks in that month. \n
- February has 28 days, but some years have 29 days. \n
- I will assume that February has 28 days, since that is the most common case. \n
- If Suzanne walks every third day, then she will walk on days 3, 6, 9, 12, 15, 18, 21, and 24. \n
- That is a total of 8 days, and she will walk 4 miles on each of those days. \n
- So, the fewest number of miles she can walk in February is 8 times 4, which is 32 miles. [Answer: 32]

Figure 7: Problem #128: Logical Reasoning Error.

822

823  
824  
825  
826  
827  
828  
829  
830  
831  
832  
833  
834  
835

## B Cross Validation

### B.1 Conflict Judges

The case #342 shows how we check the Error type of conflicting cases with the help of Judge models and human for double-check.

We use five models cross-validation to determine the type of model reasoning errors after quantization. Specifically, we use Deepseek-R1, Deepseek-V3, ChatGPT, GPT-4o and Qwen-Max as judge models, and use five models to determine the cause of the error in the reasoning process, and use the majority voting decision method to finally determine the error type.

**Problem #342:**

What integer  $x$  satisfies  $\frac{1}{4} < \frac{x}{7} < \frac{1}{3}$ ? [Gold Answer: 2]

**Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQModel):**

- To solve this inequality, I need to isolate  $x$  on one side.  $\backslash n$
- I can start by multiplying both sides by 4, since 4 is a common denominator for the fractions.  $\backslash n$
- This gives me  $\frac{4x}{7} < \frac{4}{3}$ .  $\backslash n$
- Now I can multiply both sides by 7, since 7 is a common denominator for the fractions.  $\backslash n$
- This gives me  $4x < \frac{28}{3}$ .  $\backslash n$
- To make the inequality easier to read, I can divide both sides by 4.  $\backslash n$
- This gives me  $\frac{7}{4}x < \frac{7}{1}$ .  $\backslash n$
- I can simplify the fractions by writing them as mixed numbers or decimals.  $\backslash n$
- This gives me  $1.75x < 7$ .  $\backslash n$
- Since  $x$  is an integer, I need to find the integers that satisfy this inequality.  $\backslash n$
- The smallest integer that is greater than 1.75 is 2.  $\backslash n$
- The largest integer that is less than 7 is 6.  $\backslash n$
- Therefore, the only integers that satisfy the inequality are 2, 3, 4, 5, and 6. [Answer: 2, 3, 4, 5, 6]

**Judgement of Deepseek-R1 (Primary Expert Model):**

When multiplying  $4/3$  by 7, the correct result is  $28/3 \approx 9.33$ . In step 5, the quantized model incorrectly treats  $7*(4/3)$  as  $28/3$  but then fails to correctly divide this by 4 in step 7, leading to an upper bound of 7 instead of  $7/3 \approx 2.33$ . [Execution\_Errors: computational\_error]

**Judgement of Deepseek-V3:**

The quantized model incorrectly chooses to multiply both sides by 4 instead of 7, which is the correct common denominator for the fractions. This leads to a procedural error in the subsequent steps. [Method\_Errors: Procedural error]

**Judgement of ChatGPT:**

The model incorrectly simplifies the inequality to  $\frac{7}{4}x < \frac{7}{1}$ , overlooking the upper limit should be  $\frac{28}{3}$  divided by 4, which is  $\frac{7}{3}$ , not  $\frac{7}{1}$ . [Conceptual\_Errors: contextual\_oversight]

**Judgement of GPT-4O:**

The quantized model incorrectly simplifies the range as  $\frac{7}{4}x < \frac{7}{1}$  instead of the correct range  $\frac{7}{4}x < \frac{7}{3}$ , which results in an incorrect set of integer solutions. [Method\_Errors: Procedural error]

**Judgement of Qwen-Max:**

The quantized model incorrectly simplifies the inequality  $\frac{7}{4}x < x < \frac{28}{3}$  to  $1.75 < x < 7$ . The upper bound  $\frac{28}{3}$  is approximately 9.33, not 7. [Conceptual\_Errors: conceptual\_misunderstanding]

**Human Verification Comment:**

From Step 6 to Step 7, the quantized model commits an error by incorrectly dividing both sides of the equation by 4. This aligns with the definition, which involves mistakes in arithmetic or algebraic manipulation. [Execution\_Errors: computational\_error]

Figure 8: Problem #342: Conflict Judges.

836  
837  
838  
839  
840  
841  
842  
843  
844  
845  
846  
847  
848  
849  
850  
851  
852  
853  
854  
855  
856

At the same time, we use manual intervention verification to randomly select about 2% of conflicting cases, manually determine the error type and compare it with the judgment of the judge model to ensure that the judge's vote is at a relatively correct level.

### B.2 No Errors

Although we select the extraction scripts carefully, there are still few samples where the model's output results are difficult to be correctly identified, resulting in special "No Errors" situations. This situation can usually be identified by the excellent Judge model, such as case #3812 above. The Deepseek-R1 model can reply **No Errors** after careful comparison, and we will also correct these cases among results.

Examples of easy confusion often appear in the choice of representation of decimals and fractions, including the expression between natural syntax and programming syntax, e.g.,  $1/2$  vs  $0.5$  vs  $\frac{1}{2}$  vs  $5E-01$  vs  $5 \times 10^{-1}$ .

**Problem #3812:**

What is the probability that when a fair 6-sided die is rolled, a prime number faces up? [Gold Answer:  $\frac{1}{2}$ ]

**Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQModel):**

- To find the probability of rolling a prime number on a fair 6-sided die, I need to know how many prime numbers there are on a 6-sided die and how many of them are on a fair 6-sided die.  $\backslash n$
- A prime number is a positive integer that is divisible only by 1 and itself. The prime numbers on a 6-sided die are 2, 3, and 5.  $\backslash n$
- There are 3 prime numbers on a 6-sided die, so the probability of rolling a prime number is 3 out of 6, or  $1/2$ .  $\backslash n$
- To convert this fraction to a decimal or a percentage, I can divide 1 by 2, which gives me 0.5.  $\backslash n$
- Alternatively, I can multiply 1 by 100 and divide by 2, which gives me 50.  $\backslash n$
- So the probability of rolling a prime number on a fair 6-sided die is 0.5 or 50%. [Answer: 0.5]

**Judgement of Deepseek-R1 (Primary Expert Model):**

The quantized model's answer (0.5) is numerically equivalent to the correct answer ( $1/2$ ). All reasoning steps are logically valid and mathematically correct, with no conceptual, procedural, or computational errors detected. [ No\_Error: ]

Figure 9: Problem #3812: No Error.