

000 QUANTIZATION MEETS REASONING: EXPLORING AND 001 MITIGATING DEGRADATION OF LOW-BIT LLMs IN 002 MATHEMATICAL REASONING 003

004 **Anonymous authors**
 005

006 Paper under double-blind review
 007

008 ABSTRACT 009

010 Low-bit post-training quantization (PTQ) is a practical route to deploy reasoning-
 011 capable LLMs under tight memory and latency budgets, yet it can markedly impair
 012 mathematical reasoning (drops up to 69.81% in our harder settings). We address
 013 two deployment-critical questions with process-level precision: **Where** along a
 014 step-structured solution does degradation first arise? **How** to mitigate it while
 015 staying in the low-bit regime? Across most widely used on computationally
 016 constrained scenarios PTQ methods (AWQ, GPTQ, SmoothQuant), open-source
 017 model families (Qwen, LLaMA; 0.5–7B), and math reasoning related benchmarks
 018 (GSM8K, MATH, AIME), we perform format-aligned chain-of-thought with step-
 019 aligned attribution and uncover two robust regularities: (i) PTQ disproportionately
 020 elevates method and execution errors relative to high-level conceptual mistakes; and
 021 (ii) failures emerge *early*, with the first vulnerable step flipping and cascading to
 022 the final answer. These regularities suggest a general intervention principle: restore
 023 local token-level margins exactly at the earliest failure frontier. We instantiate this
 024 principle as a lightweight *measure*→*locate*→*restore* loop that operates directly
 025 on the quantized model: detect the first faulty step, construct our "**Silver Bullet**"
 026 datasets, and apply small-scale supervised/preference tuning. In our settings, as
 027 few as 332 curated examples and 3–5 minutes of compute on a single GPU recover
 028 4-bit weight math reasoning toward the full-precision baseline while preserving
 029 PTQ efficiency. Our framework is quantizer- and architecture-agnostic within the
 030 evaluated regimes, and turns low-bit degradation from a global accuracy problem
 031 into a local, reproducible process intervention.
 032

033 1 INTRODUCTION 034

035 Transformer-based large language models (LLMs) such as LLaMA (Grattafiori et al., 2024),
 036 GPT (Achiam et al., 2023), and Qwen (Yang et al., 2024) have demonstrated strong performance on
 037 complex reasoning tasks, including mathematical competitions (Maxwell-Jia, 2025), code genera-
 038 tion (Chen et al., 2021), and logical inference (Pan et al., 2023). Yet attaining reliable accuracy on
 039 such tasks typically requires large parameter counts. The resulting inference latency and memory
 040 footprint make deploying full-precision, ultra-large models impractical in many resource-constrained
 041 scenarios. To balance resource use and accuracy, model compression has been extensively studied,
 042 including quantization (Yang et al., 2019; Rokh et al., 2023), knowledge distillation (Hinton et al.,
 043 2015; Gou et al., 2021), and pruning (Han et al., 2015). Among these, post-training quantization
 044 (PTQ) (Banner et al., 2019) lowers precision to reduce memory and improve throughput, especially
 045 on edge hardware. However, recent evidence indicates that low-bit operation (e.g., INT4) can substan-
 046 tially degrade mathematical reasoning (Feng et al., 2024; Liu et al., 2025). This raises two practical
 047 questions for deployment: **Where** does degradation emerge in the reasoning process, and **How** can it
 048 be mitigated while remaining in the low-bit regime?
 049

050 We study these questions through a systematic exploration of PTQ on widely used open-source
 051 model families and benchmarks. Concretely, we evaluate AWQ, GPTQ, and SmoothQuant on
 052 Qwen2.5 and LLaMA-3 across GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021), and
 053 AIME (Maxwell-Jia, 2025). Using format-aligned chain-of-thought and *step-aligned* attribution, we
 054 characterize quantization-induced failures across model scales and task difficulty. Two patterns are

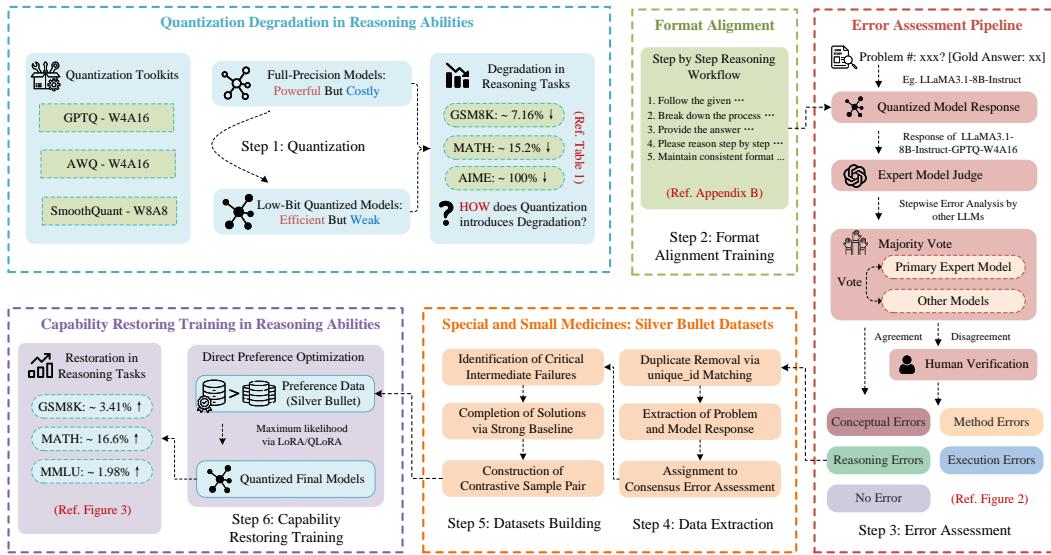


Figure 1: Pipeline of our study for investigating and restoring mathematical reasoning capabilities in quantized language models. We begin by identifying performance degradation caused by quantization, then apply format alignment training and a structured error assessment pipeline involving expert model judgments. Through this process, we analyze reasoning failures in step-by-step outputs. Targeted "Silver Bullet" datasets are constructed based on consensus error types, and used in DPO training to recover reasoning performance while maintaining the efficiency of low-bit models.

consistent: (i) PTQ predominantly increases *method* and *execution* errors (e.g., algorithm choice, rule application, carry/borrow, division/rounding), rather than high-level conceptual mistakes; and (ii) errors tend to emerge *early*, with the first vulnerable step flipping and cascading to the final answer. This diagnosis turns degradation into a targeted objective: restore token-level margins where collapse happens first.

Guided by this view, as shown on Figure 1, we adopt a lightweight *measure-locate-restore* loop that operates directly on the quantized model. We first locate the initial erroneous step, then apply small-scale supervised/preference tuning on a compact "Silver Bullet" set designed to target the observed weaknesses. In our experiments, fine-tuning on as few as 332 curated examples for 3–5 minutes on a single GPU is sufficient to recover the mathematical reasoning accuracy of W4A16 models toward their full-precision baselines, while preserving PTQ's memory and latency benefits.

We frame the study in the regime most relevant to practice: PTQ rather than quantization-aware training, so as to preserve the efficiency budget and expose unmodified low-bit failure modes. To cover current practice, we include AWQ, GPTQ, and SmoothQuant, which together span weight-only and weight-activation designs. Experiments use Qwen and LLaMA models at 0.5–7B—scales commonly deployed under constraints scenario and edge devices. Although our evidence is drawn from these choices, both the step-aligned measurement and the proposed *measure-locate-restore* loop intervention are architecture- and quantizer-agnostic by construction. Specifically, our primary contributions are as follows:

- We build a step-aligned measurement suite and hierarchical error taxonomy that expose a robust PTQ-induced shift toward *method* and *execution* errors, with earlier first-step flips, consistent across the most popular models, bit-widths, and benchmarks.
- We develop an automated chain-of-thought error-analysis pipeline (judge ensemble and light human audit) that attains 97.2% labeling accuracy on 9,908 failure cases, enabling fine-grained, reproducible attribution by error type and first faulty step.
- We introduce a compact *measure-locate-restore* loop that tunes the quantized model with targeted "Silver Bullet" pairs, recovering W4A16 mathematical reasoning to near full precision with just 332 curated examples and 3–5 minutes on a single GPU—without access to pretraining data.

108

2 RELATED WORKS

109

2.1 QUANTIZATION METHODS

111 Quantization is a computational efficiency optimization technique that maps high-precision tensors
 112 $X \in \mathbb{R}^{m \times n}$ into low-bit discrete representations. This work focuses on hardware-efficient uniform
 113 quantization, a linear mapping paradigm particularly suited for deployment on embedded systems
 114 with fixed-point arithmetic units. The method achieves significant reductions in model storage
 115 requirements and inference energy consumption while maintaining computational tractability.

116 For b -bit quantization, the mathematical formulation is expressed as:

$$118 \quad \hat{X} = Q(X; b) = s \cdot \Pi_{\Omega(b)} \left(\frac{X}{s} \right) \quad (1)$$

121 where the quantization step size $s = \frac{\max(X) - \min(X)}{2^b - 1}$ dynamically adapts to the input distribution,
 122 effectively compressing the continuous floating-point space into an integer set $\Omega(b) = \{0, 1, \dots, 2^b - 1\}$. The projection function $\Pi(\cdot)$ discretizes normalized values through nearest-neighbor rounding,
 123 with the rounding error being a primary source of quantization-induced precision loss. Notably, the
 124 step size s governs the resolution of quantization intervals—larger dynamic ranges may sacrifice
 125 fine-grained details, necessitating calibration strategies for optimal parameter selection in practical
 126 implementations.

128 The engineering trade-offs of quantization manifest in multiple dimensions:

- 130 • **Bit-width Flexibility:** While aggressive 4-bit quantization reduces model size to 1/8 of its
 131 original footprint, it risks substantial accuracy degradation. Conversely, 8-bit quantization
 132 typically achieves near-full-precision performance in most scenarios.
- 133 • **Dynamic vs. Static Modes:** Dynamic quantization computes step sizes at runtime to adapt
 134 to input variations, whereas static quantization pre-calibrates parameters offline to minimize
 135 inference overhead.
- 136 • **Weight-only vs. Weight-activation:** Weight-only quantization restricts low-bit representa-
 137 tion to model parameters, preserving activation precision for tasks sensitive to numerical
 138 stability. In contrast, weight-activation quantization jointly compresses both weights and
 139 intermediate activations, achieving higher memory efficiency at the cost of error accumula-
 140 tion.

141 Our methodology encompasses two complementary quantization approaches: (1) Post-training
 142 weight-only compression via AWQ (Lin et al., 2024) and GPTQ (Frantar et al., 2022), achieving 4-bit
 143 precision preservation through adaptive rounding strategies; (2) The SmoothQuant (Xiao et al., 2023)
 144 framework for joint weight-activation quantization, maintaining 8-bit numerical stability via learned
 145 scale migration. This dual-strategy design addresses distinct precision requirements: aggressive
 146 weight compression for memory efficiency versus moderate activation quantization for computational
 147 robustness. Comprehensive implementation protocols, including gradient-aware quantization grid
 148 adaptation and layer-wise sensitivity analysis, are detailed in Appendix A.

149

2.2 REASONING ABILITY OPTIMIZATION IN LARGE LANGUAGE MODELS

151 LLMs increasingly demonstrate strong general-purpose reasoning skills, spanning commonsense
 152 inference to domain-specific problem solving. Early evidence from Minerva (Lewkowycz et al., 2022)
 153 shows that scaling models and tailoring data can unlock advanced mathematical competence—one
 154 instance of the broader trend that rich intermediate computations boost reasoning fidelity. Prompt-
 155 engineering techniques such as Chain-of-Thought (Wei et al., 2022) and its code-generating variant
 156 Program-of-Thought (Chowdhery et al., 2023) further improve multi-step reasoning by encouraging
 157 models to decompose tasks into interpretable sub-steps.

158 Orthogonal to prompting, alignment research pursues systematic post-training refinements. Instruc-
 159 tion tuning on diversified task mixtures (FLAN) (Wei et al., 2021) and lightweight data-curation
 160 pipelines (Alpaca) (Taori et al., 2023) make models broadly helpful, while Direct Preference Opti-
 161 mization (DPO) (Rafailov et al., 2024) offers sample-efficient preference learning without full RLHF.
 Reliability has been pushed along two complementary axes: self-consistency voting for answer

162 selection (Wang et al., 2022) and process-level supervision with stepwise reward models (Lightman
 163 et al., 2023b), both grounded in verifiable-reasoning theory (Creswell & Shanahan, 2022).

164
 165 Building on these insights, we adopt process-supervised fine-tuning that obliges the model to articulate
 166 and justify each intermediate step. This explicit trace makes it possible to localize—and later
 167 ameliorate—reasoning failures introduced by low-bit quantization, providing a principled path toward
 168 efficient yet reliable LLM deployment.

169 3 METHODOLOGY

170 3.1 QUANTIZATION-INDUCED DEGRADATION: MEASUREMENT AND ATTRIBUTION

171
 172 In this section, we investigate how low-bit quantization influences the reasoning performance of LLMs.
 173 Distinct from prior works, we examine each model’s step-by-step solution trajectory and conduct a
 174 fine-grained quantitative–qualitative error analysis to pinpoint the root causes of reasoning failures.
 175 Our study centers on mathematically oriented tasks, which serve as a rigorous and representative
 176 proxy for general reasoning ability.

177 3.1.1 QUANTIZATION

178
 179 We conduct a comprehensive investigation into the effects of quantization techniques, encompasses
 180 two complementary quantization approaches: (1) Post-training weight-only compression via AWQ
 181 (Lin et al., 2024) and GPTQ (Frantar et al., 2022), achieving 4-bit weight precision preservation
 182 through adaptive rounding strategies and keep the data format of activations in 16-bit; (2) The
 183 SmoothQuant (Xiao et al., 2023) framework for joint weight-activation quantization, maintaining
 184 8-bit numerical stability via learned scale migration. Through the systematic application of these
 185 most popular and wild-use quantization techniques, we provide a rigorous and balanced analysis of
 186 the resulting quantized models, offering valuable insights into their performance characteristics and
 187 trade-offs. Detailed algorithmic descriptions and mathematical derivations for all three methods are
 188 provided in Appendix A.

189 3.1.2 FORMAT ALIGNMENT TRAINING

190
 191 To address the challenge of inconsistent instruction following and irregular output formatting in
 192 model-generated solutions, we introduce a **format alignment stage**. This phase aims to instill in the
 193 model a structured, step-by-step reasoning workflow without altering its underlying mathematical
 194 knowledge. Crucially, the objective here is **NOT** to teach the model new mathematical facts or
 195 knowledge injection, but rather to ensure strict adherence to a prescribed output format, thereby
 196 enabling reliable qualitative and quantitative analysis of reasoning capability across quantized and
 197 full-precision variants.

198
 199 We employ LoRA (Hu et al., 2021) and QLoRA (Dettmers et al., 2024) for full-precision model
 200 and quantized model respectively as lightweight adaptation techniques for format alignment. These
 201 methods efficiently align knowledge of step-by-step solution formats into the model’s latent space
 202 without extensive retraining. This fine-tuning enables us to observe how multi-step reasoning is
 203 preserved or altered once the model is quantized, offering deeper insights into any capability loss
 204 induced by compression.

205
 206 For alignment, we utilize the PRM800K dataset (Lightman et al., 2023a), which provides 800K step-
 207 level correctness annotations from 75K solutions to 12K problems. These annotations supply granular,
 208 step-by-step reasoning trajectories, equipping models to separate complex problem-solving processes
 209 into well-defined stages. To reinforce this structure, we adopt a consistent system prompt across
 210 training and evaluation, ensuring that the boundaries of logical steps and final answers are clearly
 211 delineated. This consistent, step-by-step alignment is a necessary foundation for our subsequent
 212 qualitative and quantitative analyses of quantization-induced degradation in mathematical reasoning.
 213 More details are presented on Appendix B

214 3.1.3 DETAILED EXAMINATION OF REASONING PROCESS

215
Qualitative Analysis. To systematically investigate the underlying reasons for degradation in
 216 quantized models, we performed a qualitative error analysis inspired by established categorizations

216 Table 1: Comparison of quantization methods applied to the Llama-3 and Qwen2.5 model families.
 217 AWQ and GPTQ employ 4-bit weight and 16-bit activation quantization, whereas **SQ (SmoothQuant)**
 218 uses 8-bit weight and 8-bit activation quantization.

GSM8K			MATH			AIME							
	Van.	AWQ (W4A16)	GPTQ (W4A16)	SQ (W8A8)	Van.	AWQ (W4A16)	GPTQ (W4A16)	SQ (W8A8)	Van.	AWQ (W4A16)	GPTQ (W4A16)	SQ (W8A8)	
Llama	3.1-8B-Inst.	79.98	79.53 (0.56%)	78.85 (1.41%)	80.14 (-0.20%)	50.72	46.44 (8.44%)	46.04 (9.23%)	50.26 (0.91%)	10	1.11 (88.90%)	5.56 (44.40%)	6.67 (33.30%)
	3.2-3B-Inst.	77.26	74.45 (3.64%)	71.57 (7.36%)	77.71 (-0.58%)	45.82	40.76 (11.04%)	42.66 (6.90%)	45.74 (0.17%)	4.44	0 (100%)	2.22 (50%)	4.44 (0)
	3.2-1B-Inst.	45.56	39.12 (14.14%)	39.58 (13.13%)	45.19 (0.81%)	20.58	16.2 (21.28%)	15.78 (23.32%)	20.96 (-1.85%)	3.33	0 (100%)	0 (100%)	0 (100%)
	7B-Inst.	87.04	86.5 (0.62%)	85.14 (2.18%)	86.73 (0.36%)	72.48	69.84 (3.64%)	69.6 (3.97%)	72.04 (0.61%)	11.11	10 (9.99%)	8.89 (19.98%)	8.89 (19.98%)
	3B-Inst.	81.35	79.68 (2.05%)	79.76 (1.95%)	81.27 (0.1%)	63.3	56 (11.53%)	55.02 (13.08%)	63.52 (-0.35%)	4.44	3.33 (25%)	3.33 (25%)	2.22 (50%)
	1.5B-Inst.	68.23	61.11 (10.44%)	59.89 (12.22%)	68.46 (-0.34%)	43.74	25.6 (41.47%)	31.2 (28.67%)	43.52 (0.5%)	2.22	1.11 (50%)	0 (100%)	2.22 (0)
Qwen2.5	0.5B-Inst.	43.37	27.07 (37.58%)	26.61 (38.64%)	41.55 (4.2%)	23.98	8.02 (66.56%)	7.24 (69.81%)	24 (-0.08%)	0	0 (0)	0 (0)	0 (0)

231 *Notes: Van. denotes the vanilla full-precision baseline; Inst. is an abbreviation of Instruct.*

233 from previous literature (Brown et al., 2016), (Delastri & Lolang, 2023) and (Kurudirek et al., 2023),
 234 which categorize **real world student errors** in mathematical problem solving. Building on these
 235 frameworks, we conduct a qualitative analysis by classifying model-generated errors into seven
 236 fine-grained error types, organized under four high-level categories. The definitions of these error
 237 types are detailed as follows:

- 238 • **Conceptual Errors** arise when the model fundamentally misunderstands the underlying
 239 principles or context. This includes misgrasping core theories or overlooking domain-
 240 specific constraints (e.g., boundary conditions), leading to distorted problem framing and
 241 invalid solutions.
- 242 • **Method Errors** occur when mathematical methods are misapplied or chosen inappropriately.
 243 Typical cases include executing standard algorithms incorrectly, skipping key procedures, or
 244 misusing formulae in unsuitable contexts.
- 245 • **Execution Errors** stem from mistakes in arithmetic or symbolic manipulation, such as
 246 faulty calculations, erroneous expansions, or mislabeling variables. These slips compromise
 247 intermediate computations and ultimately the final answer.
- 248 • **Reasoning Errors** reflect flaws in logical flow, where inference steps do not follow coher-
 249 ently or essential links are omitted, creating gaps that render the conclusion unsupported.

250 **Quantitative Analysis and Error Assessment Pipeline.** To facilitate a rigorous and scalable
 251 evaluation of quantization-induced errors in reasoning tasks, we developed an automated assessment
 252 pipeline powered by state-of-the-art language models. This pipeline systematically processes model
 253 outputs and classifies errors according to our predefined *error_types_list* taxonomy. By leveraging
 254 a pre-trained transformer as the core evaluator, we reduce subjective bias and ensure consistent,
 255 reproducible error analyses across all experimental conditions. Furthermore, the computational
 256 scoring framework supports high-throughput performance assessment while preserving granularity in
 257 error categorization.

258 Our quantitative assessment pipeline comprises three primary stages:

259 **1. Expert Model Judgement:** For each instance in which a quantized model produces an incorrect
 260 answer, we employ a dedicated "expert model" to analyze the error. The expert model is tasked
 261 with: (a) identifying the first occurrence of an error, (b) specifying the exact step where the error is
 262 introduced, (c) assigning an error category based on a nested classification scheme, and (d) providing
 263 an explanation along with a confidence score for its determination.

264 **2. Majority Voting:** To curb hallucinations and improve evaluation reliability, we apply a three-stage
 265 majority-vote protocol to the outputs of five language models—*DeepSeek-R1* (Guo et al., 2025)
 266 (primary), GPT-4o, GPT-4, Qwen-Max (Yang et al., 2025), and *DeepSeek-V3* (Liu et al., 2024).
 267 Instances of disagreement are flagged for further review, ensuring consistency and minimizing
 268 spurious judgments. **Rule1-Unanimous agreement:** If all four auxiliary models concur with the
 269 reference judgment from *DeepSeek-R1*, the answer is accepted. **Rule2-Simple majority:** If exactly

270 three auxiliary models concur with DeepSeek-R1, the answer is likewise accepted. **Rule3-Escalation:**
 271 Otherwise, the instance is forwarded to two independent human annotators for arbitration.
 272

273 **3. Human Annotation:** For cases with conflicting assessments from the majority vote, we introduce
 274 two human annotators to manually review is conducted. The annotator need to follow the anno-
 275 tation document and review the explanations of five expert models then give the final assessment.
 276 Additionally, we also randomly sample 2% of the passed evaluated cases to verify the accuracy and
 277 consistency of the automated judgments. The annotation documents are detailed in Appendix C.

278 This pipeline is intentionally designed to be conservative and to avoid spurious "false consensus"
 279 in the automatic labels. Among several strong candidate judges, we select DeepSeek-R1 as the
 280 pivot because its reflective, stepwise chain-of-thought makes it particularly suitable for localizing the
 281 first erroneous step and producing structured explanations; in a small pilot on a random subset of
 282 failures, its error-type predictions also showed the highest agreement with two human annotators. For
 283 each incorrect model output, we then collect judgments from all five expert models and accept an
 284 automatic label only when at least three of the four auxiliary models concur with the pivot, that is, at
 285 least four out of five models agree on both the error type and its location; otherwise the instance is
 286 escalated to human annotators. In addition, we randomly sample two percent of the automatically
 287 accepted cases for manual audit.

288 Under this protocol, our automated error-assessment pipeline matches the final human judgment on
 289 97.2% of 9,908 failure cases, with the remaining discrepancies concentrated in borderline situations,
 290 for example when the canonical answer is " $\frac{11}{2}$ " but the quantized model outputs the
 291 numerically equivalent "5.5" and some expert models still flag an error due to subtle reasoning or
 292 formatting differences. These observations suggest that our judge framework achieves high precision
 293 at the cost of slightly lower recall, which is appropriate for the downstream analyses in this paper;
 294 further implementation details and annotation guidelines are provided in Appendix C.

295 3.2 RESTORING REASONING ABILITIES IN QUANTIZED MODELS

296 3.2.1 DATA EXTRACTION

297 Building on the analysis in Section 3.1.3, we construct our evaluation subset by filtering and cate-
 298 gorizing problem instances according to model error types. First, to eliminate any risk of data leakage,
 299 we remove all overlapping examples between the MATH and MATH-500 test sets by matching on
 300 their `unique_id` fields. Next, for each quantized model, we identify those problems that the full-
 301 precision counterpart answers correctly but on which the quantized variant fails, based on the models'
 302 majority-vote outputs. We then collect the corresponding problem prompts and model-generated
 303 responses for these failure cases. Finally, leveraging the labels produced by our error-assessment
 304 pipeline, we assign each case to its consensus error category for downstream analysis.

305 3.2.2 SILVER BULLET DATASETS BUILDING

306 During the execution of our error-assessment pipeline, we identify and record the exact reasoning
 307 step at which each quantized model initially commits an error. Our qualitative analysis indicates that
 308 many reasoning failures originate from incorrect intermediate computations or boundary adjustments,
 309 on which all subsequent solution steps heavily depend. Leveraging this observation, we construct
 310 a targeted counterexample dataset by truncating the incorrect reasoning traces precisely at the
 311 identified erroneous steps. Subsequently, we prompt powerful baseline models (Llama-3.2-70B
 312 and Qwen2.5-Max) to resume and complete these truncated solutions until the correct answers are
 313 derived. Consequently, we designate the original quantized models' erroneous partial solutions as
 314 negative samples, while adopting the accurately completed solutions generated by the larger models as
 315 positive samples. This approach yields our "Silver Bullet" datasets, specifically designed to facilitate
 316 downstream error correction and model fine-tuning.

317 3.2.3 CAPABILITY RESTORING TRAINING

318 To reclaim the reasoning capability lost after low-bit quantization, we fine-tune each *quantized* model
 319 using **Direct Preference Optimization** (DPO) (Rafailov et al., 2024). Given a prompt x and a pair
 320 of responses (y^+, y^-) where y^+ is the correct answer and y^- is the quantized model's incorrect

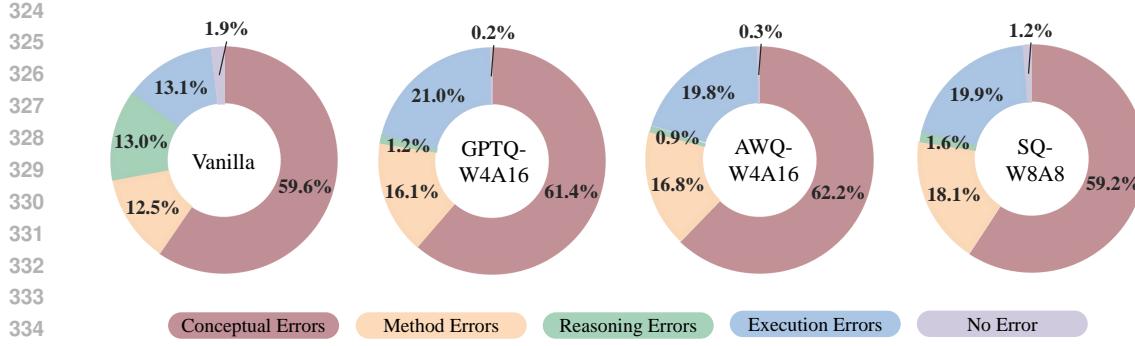


Figure 2: Error assessment results for full-precision and quantized models. For the full-precision model, we aggregate all problems it answered incorrectly; for each quantized model, we count only those problems that the full-precision model solved correctly but the quantized model failed, enabling comparison of quantization-induced changes across error dimensions.

answer, y^+ is preferred to y^- , DPO maximizes the log-likelihood gap between the two while softly constraining the new policy π_θ toward the frozen reference policy π_{ref} . The objective is

$$\mathcal{L}_{\text{DPO}}(\theta) = \mathbb{E}_{(x, y^+, y^-) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left[\log \pi_\theta(y^+ | x) - \log \pi_\theta(y^- | x) \right. \right. \right. \\ \left. \left. \left. - (\log \pi_{\text{ref}}(y^+ | x) - \log \pi_{\text{ref}}(y^- | x)) \right] \right) \right]. \quad (2)$$

where σ is the sigmoid function and β is an inverse-temperature hyper-parameter (we set $\beta = 1$). Because the reference gap is constant with respect to θ , maximizing \mathcal{L}_{DPO} is equivalent to minimizing $\text{KL}(\pi_\theta \parallel \pi_{\text{ref}})$ subject to pairwise preference constraints, thus yielding a stable, RL-free preference-alignment procedure with solid theoretical footing.

We realize the adaptation using LoRA and 4-bit QLoRA. Across all experiments, we set the LoRA rank to 32 for every injected adapter matrix and optimize with a cosine learning-rate schedule (base learning rate 1×10^{-6} , warm-up ratio 0.1) under a global batch size of 8. Training minimizes the sigmoid preference loss implied by \mathcal{L}_{DPO} .

4 EXPERIMENTS

4.1 EVALUATING QUANTIZATION EFFECTS

In this phase of our study, we selected three benchmark datasets of varying difficulty levels to evaluate the degradation introduced by quantization across different reasoning complexities.

- **GSM8K** is a high-quality dataset of grade-school level math word problems released by OpenAI, containing 8,500 problems that typically require 2 to 8 steps of reasoning.
- **MATH** is a more challenging dataset composed of 12,500 competition-level high school math problems, covering seven mathematical domains including algebra, geometry, number theory, and probability and statistics, generally requires 15 or more steps of logical reasoning.
- **AIME** (American Invitational Mathematics Examination) is a high-difficulty International Mathematical Olympiad(IMO) competition designed for advanced middle and high school students with 90 problems (we combine problems from 2022-2025 for a widely evaluation).

We maintain consistency in both the global batch size and the prompt with those used during alignment and evaluation. This setup ensures a fair comparison across all models. According to the Table 1 we find these two trends:

Smaller-scale models suffer more severe losses in complex reasoning ability after quantization: Across all quantization methods, smaller-scale models consistently demonstrate increased vulnerability to quantization. Specifically, the Qwen2.5-0.5B-Instruct model experiences accuracy drops exceeding 60% post-quantization, whereas the larger Qwen2.5-7B-Instruct model incurs only a

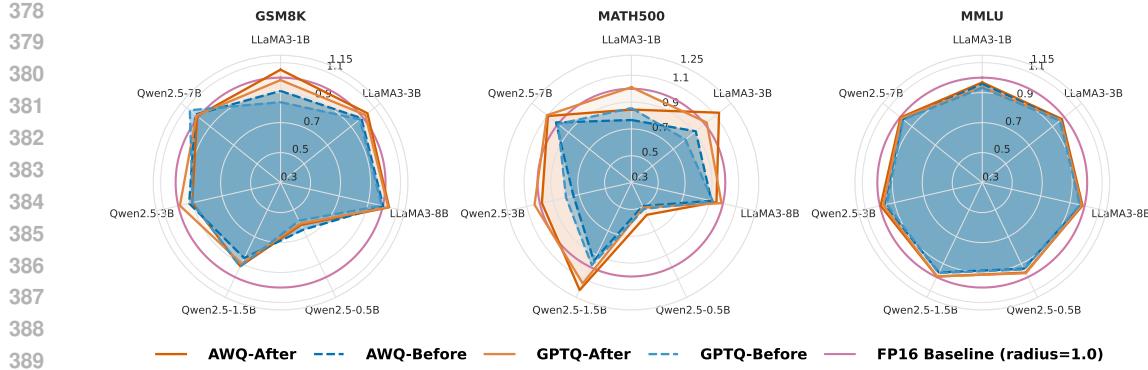


Figure 3: **Relative capability restoration with our method.** Radar values are normalized to each model’s Vanilla-FP16 accuracy on the same benchmark (radius 1.0). Solid = After Restoration, dashed = Before Restoration (AWQ, GPTQ).

modest degradation of approximately 2–3%. This trend is also corroborated within the Llama3 model series. To rule out potential biases arising from larger models more readily fitting the calibration datasets, we further validated our findings using calibration datasets of varying sizes, consistently obtaining similar results. This evidence suggests that smaller models are more adversely affected by quantization-induced shifts in feature distributions, thereby experiencing more severe performance declines in complex mathematical reasoning tasks.

Performance degradation becomes more pronounced as the task complexity increases: We evaluated model accuracy across three mathematical reasoning benchmarks of varying difficulty levels. Our results indicate a clear trend wherein performance degradation exacerbates as task complexity rises. Among these, AIME represents the most challenging benchmark, with even full-precision models constrained by their scale unable to solve all problems effectively. The MATH dataset, characterized by evenly distributed difficulty tiers, poses intermediate-level complexity, while GSM8K is comparatively less challenging. Notably, quantized models exhibited relatively minor accuracy losses on the simpler GSM8K benchmark, with an average performance decline of only 7.16%. In contrast, the MATH dataset incurred a more pronounced average degradation of 15.18%. The most severe impact was observed on the highly challenging AIME benchmark, where quantization frequently led to complete failure in problem-solving capability.

We also evaluate the quantization-induced degradation on both thinking-mode models and larger-scale models for mathematical reasoning tasks, the detailed results are reported in Appendix E.1. Together with the corresponding evaluations under the same quantization settings on general-purpose benchmarks, the detailed results are reported in Appendix E.2.

4.2 ERROR TAXONOMY AND ITS SHIFT UNDER QUANTIZATION

Error profile of full-precision models. Using the assessment pipeline in Section 3.1.3, we examined every problem that the full-precision models answered incorrectly. *Conceptual Errors* were the most frequent (59.6%), while *Method*, *Execution*, and *Reasoning* Errors appeared at comparable rates of 12.5%, 13.0%, and 13.1%, respectively.

Impact of quantization. We next analyzed the subset of problems that full-precision models answered correctly but failed under quantized models. Across all three quantization methods, we observed a noticeable increase in the proportion of *Method Errors* and *Execution Errors*, suggesting that quantization predominantly impairs the model’s ability to perform procedural operations and arithmetic execution (Feng et al., 2024). Supporting this observation, our case study reveals that quantized models exhibit greater difficulty in handling tasks involving basic arithmetic operations and numerical computation.

Why Reasoning Errors seem to vanish. The apparent reduction in *Reasoning Errors* after quantization arises from a statistical masking effect induced by our standardized evaluation protocol. Each trajectory is assigned a single error type according to the first erroneous step. In contrast, reasoning type failures in our taxonomy usually occur later in the solution when global logical consistency or boundary conditions are evaluated. Quantization increases the likelihood of earlier and simpler mistakes such as *Conceptual* or *Execution Errors*, and these early failures “hide” subsequent reasoning

Table 2: Ablation results on GSM8K, MATH500, and MMLU under AWQ/GPTQ (W4A16). Row labels denote training subsets: ALL-S = all error cases with step-aligned supervision from the first error; CE/ME/EE-S = only conceptual/method/execution errors with step alignment; Rand-NS = size-matched random sampling without step alignment; ALL-NS = all error cases without step alignment. Numbers are accuracy (%).

	Llama-3.2-3B-Inst						Qwen2.5-3B-Inst.						Avg.	
	AWQ(W4A16)			GPTQ(W4A16)			AWQ(W4A16)			GPTQ(W4A16)				
	GSM8K	MATH 500	MMLU	GSM8K	MATH 500	MMLU	GSM8K	MATH 500	MMLU	GSM8K	MATH 500	MMLU		
ALL-S	74.3	36.8	60.57	73.01	33	59.9	68.84	38.4	64.8	75.21	40.6	63.63	57.42	
CE-S	73.19	35.4	60.42	73.31	31.6	59.98	70.28	35.6	65	75.28	34.6	63.61	56.52	
ME-S	73.62	32	60.45	72.4	30.4	59.86	69.6	34.8	65.01	76.27	32.6	63.81	55.90	
EE-S	73.39	31.2	60.47	72.63	31.2	59.95	69.29	34.8	64.95	76.12	34	63.84	55.98	
Rand-NS	73.84	30.9	60.51	72.71	31.2	59.97	70.05	32.5	64.92	74.13	34.4	63.71	55.73	
ALL-NS	70.74	27.8	60.25	69.45	15	59.94	66.79	29.4	65.01	71.04	22.4	64.1	51.83	

flaws from the statistics. Our case study confirms that many trajectories labeled as other error types still contain additional reasoning problems at later steps, although these later issues are not recorded because they occur after the first mistake. This masking effect therefore reflects how quantization reshapes the distribution of observed first errors under a reproducible and unambiguous protocol, not an actual disappearance of deeper reasoning mistakes. Absolute accuracies and examples are provided in Appendix D.

4.3 CAPABILITY RESTORATION

To prevent data leakage during evaluation, we report results on **MATH-500** (Lightman et al., 2023b), a 500-problem set that is disjoint from PRM800K yet mirrors the original MATH benchmark in topic coverage and difficulty. Performance on MATH-500 thus reflects genuine reasoning recovery rather than memorization. We also measure accuracy on GSM8K and MMLU (Hendrycks et al., 2020) to assess how well the restored model generalises to other reasoning-intensive tasks. The results are visually presented in Figure 3, with additional details provided in Appendix F.

4.4 ABLATION STUDY

To isolate the contributions of each component in our quantization recovery pipeline, we perform a series of ablation studies. Unless otherwise noted, all runs fix the training budget, optimizer, prompts, and decoding policy. We compare four variants of our training pairs (the *failure subset* refers to instances the quantized model answers incorrectly under baseline evaluation):

- ALL-STEP: all error cases from the failure subset; step-aligned supervision *resumes at the first-error step* (i.e., the first step where the model deviates from the gold solution) and continues step-wise along the gold trajectory to the final answer.
- CE/ME/EE-STEP: identical to ALL ERRORS-STEP but restricted to a single error type (*Conceptual-Error / Method-Error / Execution-Error*), with supervision resuming from the first-error step.
- RANDOM-NONSTEP: size-matched random sampling from the math corpus, independent of whether the quantized model fails; positives are the gold solutions, and *no* first-error resuming is applied.
- ALL-NONSTEP: all error cases from the failure subset, but *without* resuming from the first-error step; positives are the full gold solutions.

4.5 DISCUSSION

Synthesizing the results from Sections 4.3 and 4.4 together with the trends in Figure 2 and Table 2, we draw three main conclusions:

(i) Targeted recovery with our "Silver Bullet" datasets. Fine-tuning on the compact failure-targeted split substantially restores performance on MATH500 while also boosting GSM8K, and does so without hurting broad-domain reasoning as measured by MMLU. This intervention uses only a few hundred preference pairs and a few minutes of training on a single GPU, yet closes much of the gap between quantized and full-precision models. This confirms that the "*Silver Bullet*" provides *sample-efficient capability recovery rather than simple memorization*.

486 (ii) **Quantization disproportionately erodes procedural and executional skills.** Our error taxonomy
 487 shows that weight–activation quantization mainly increases *method* and *execution* errors—such
 488 as carrying out multi-step arithmetic or handling boundary conditions—rather than high-level con-
 489 ceptual reasoning. Because these mistakes often occur early, they propagate to invalidate otherwise
 490 correct derivations, explaining the steep drop on math-centric tasks.

491 (iii) **Step-wise positives outperform naive alternatives.** Ablation results in Table 2 is run under a
 492 strictly matched data and compute budget: each setting uses the same number of stepwise preference
 493 pairs and identical training hyperparameters, and only the selection of problems and traces is varied.
 494 Under this controlled setup, training on our error-targeted stepwise ALL split consistently outperforms
 495 both size-matched random supervision (RANDOM) and the non-stepwise variant (NON-STEP) that
 496 adopts full-precision derivations without restarting from the first erroneous step. On average, ALL
 497 improves accuracy by about 1.7 points over RANDOM and by about 5.6 points over NON-STEP,
 498 with the largest gap of 6.2 points on MATH500 for the most quantization-sensitive configuration
 499 (Qwen2.5-3B with GPTQ). These results indicate that locating the earliest quantization-induced
 500 failure and regenerating the remaining steps from that point provides a much stronger learning signal
 501 than unconditioned math supervision or full-solution positives, especially in the low-data regime we
 502 consider.

503 5 CONCLUSION

504 In this study, we present a systematic study of quantization-induced degradation in the mathematical
 505 reasoning abilities of large language models, revealing that low-bit post-training quantization
 506 especially harms smaller models’ procedural and execution skills. To address this, we propose a
 507 lightweight recovery pipeline that combines step-aligned error analysis with targeted fine-tuning on
 508 compact, automatically constructed “Silver Bullet” datasets. Experiments show that, with minimal
 509 data and compute, quantized models can recover reasoning performance to match their full-precision
 510 counterparts while preserving efficiency and general capabilities. Our approach offers a practical
 511 and extensible solution for deploying quantized LLMs in resource-constrained settings, and opens
 512 avenues for robust reasoning restoration in broader domains.

514 REFERENCES

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

519 Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional
 520 networks for rapid-deployment. *Advances in Neural Information Processing Systems*, 32, 2019.

521 Janice Brown, Kim Skow, and Center IRIS. Mathematics: Identifying and addressing student errors.
 522 *The Iris Center*, 31, 2016.

524 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared
 525 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large
 526 language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

528 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam
 529 Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm:
 530 Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113,
 531 2023.

532 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and
 533 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.
 534 *arXiv preprint arXiv:1803.05457*, 2018.

535 Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser,
 536 Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve
 537 math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

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

540 Lusiana Delastri and Enos Lolang. Students' conceptual error and procedural error in solving
 541 algebraic problems. *Multicultural Education*, 9(1):18–24, 2023.
 542

543 Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient fine-tuning
 544 of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024.
 545

546 Guhao Feng, Kai Yang, Yuntian Gu, Xinyue Ai, Shengjie Luo, Jiacheng Sun, Di He, Zhenguo Li,
 547 and Liwei Wang. How numerical precision affects mathematical reasoning capabilities of llms.
 548 *arXiv preprint arXiv:2410.13857*, 2024.
 549

550 Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. Gptq: Accurate post-training
 551 quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022.
 552

553 Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A
 554 survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
 555

556 Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad
 557 Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of
 558 models. *arXiv preprint arXiv:2407.21783*, 2024.
 559

560 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
 561 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
 562 via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
 563

564 Song Han, Huiyi Mao, and William J Dally. Deep compression: Compressing deep neural networks
 565 with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
 566

567 Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and
 568 Jacob Steinhardt. Measuring massive multitask language understanding. *arXiv preprint
 569 arXiv:2009.03300*, 2020.
 570

571 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,
 572 and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv
 573 preprint arXiv:2103.03874*, 2021.
 574

575 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv
 576 preprint arXiv:1503.02531*, 2015.
 577

578 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
 579 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint
 580 arXiv:2106.09685*, 2021.
 581

582 Abdullah Kurudirek, Bnar Karim, Delan Sarhang, and Safarov Tulqin. Math misconceptions:
 583 Mistakes, misunderstanding, and confusion. 2023.
 584

585 Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ra-
 586 masesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, et al. Solving quantitative
 587 reasoning problems with language models. *Advances in Neural Information Processing Systems*,
 588 35:3843–3857, 2022.
 589

590 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 591 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. *arXiv preprint
 592 arXiv:2305.20050*, 2023a.
 593

594 Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan
 595 Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step, 2023b. URL
 596 <https://arxiv.org/abs/2305.20050>.
 597

598 Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan
 599 Xiao, Xingyu Dang, Chuang Gan, and Song Han. Awq: Activation-aware weight quantization for
 600 on-device llm compression and acceleration. *Proceedings of Machine Learning and Systems*, 6:
 601 87–100, 2024.
 602

594 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao,
 595 Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*
 596 *arXiv:2412.19437*, 2024.

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

600 Maxwell-Jia. AIME_2024 — hugging face datasets. https://huggingface.co/datasets/Maxwell-Jia/AIME_2024, 2025. Accessed 2025-04-06.

601 Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. Logic-lm: Empowering
 602 large language models with symbolic solvers for faithful logical reasoning. *arXiv preprint*
 603 *arXiv:2305.12295*, 2023.

604 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 605 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances*
 606 *in Neural Information Processing Systems*, 36, 2024.

607 Babak Rokh, Ali Azarpeyvand, and Alireza Khataymoori. A comprehensive survey on model
 608 quantization for deep neural networks in image classification. *ACM Transactions on Intelligent*
 609 *Systems and Technology*, 14(6):1–50, 2023.

610 Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. Commonsenseqa: A question
 611 answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of*
 612 *the North American Chapter of the Association for Computational Linguistics: Human Language*
 613 *Technologies, Volume 1 (Long and Short Papers)*, pp. 4149–4158, 2019.

614 Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy
 615 Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.

616 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-
 617 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.
 618 *arXiv preprint arXiv:2203.11171*, 2022.

619 Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du,
 620 Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint*
 621 *arXiv:2109.01652*, 2021.

622 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 623 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*
 624 *neural information processing systems*, 35:24824–24837, 2022.

625 Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant:
 626 Accurate and efficient post-training quantization for large language models. In *International*
 627 *Conference on Machine Learning*, pp. 38087–38099. PMLR, 2023.

628 An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li,
 629 Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint*
 630 *arXiv:2412.15115*, 2024.

631 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang
 632 Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*,
 633 2025.

634 Jiwei Yang, Xu Shen, Jun Xing, Xinmei Tian, Houqiang Li, Bing Deng, Jianqiang Huang, and
 635 Xian-sheng Hua. Quantization networks. In *Proceedings of the IEEE/CVF conference on computer*
 636 *vision and pattern recognition*, pp. 7308–7316, 2019.

637 Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine
 638 really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

639 Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny
 640 Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint*
 641 *arXiv:2311.07911*, 2023.

648 A APPENDIX A
649
650651 A.1 AWQ
652653 AWQ (Activation–Aware Weight Quantization) compensates for the long-tailed distribution of
654 activations before the weight tensor is discretised. Let $\mathbf{A} \in \mathbb{R}^{B \times d}$ be the mini-batch activations and
655 $\mathbf{W} \in \mathbb{R}^{d \times m}$ the corresponding weights. A positive scale vector $\gamma \in \mathbb{R}_+^d$ is chosen such that
656

657
$$\tilde{\mathbf{Y}} = (\mathbf{A} \odot \gamma^{-1})(\gamma \odot Q(\mathbf{W}))^\top, \quad \gamma_k = (\text{mean } |\mathbf{A}_{:,k}|^\alpha)(\text{mean } |\mathbf{W}_{k,:}|^{-\beta}),$$

658

659 where $(\alpha, \beta) \in [0, 1]$ control the balance between activation and weight magnitudes and $Q(\cdot)$ denotes
660 an asymmetric 4-bit quantiser. Because the rescaling is folded back into \mathbf{W} , the forward pass is
661 identical to the unscaled INT4 kernel and incurs no extra latency.
662663 A.2 GPTQ
664665 GPTQ formulates post-training quantisation as a blockwise least-squares problem over a small
666 calibration set $\mathcal{C} = \{\mathbf{A}^{(i)}\}_{i=1}^{|\mathcal{C}|}$:
667

668
$$\tilde{\mathbf{W}} = \arg \min_{\mathbf{W}' \in \mathcal{Q}} \sum_{i=1}^{|\mathcal{C}|} \|\mathbf{W}' \mathbf{A}^{(i)} - \mathbf{W} \mathbf{A}^{(i)}\|_F^2,$$

669

670 where \mathcal{Q} is the set of weight tensors representable by the target bit-width. The optimisation pro-
671 ceeds greedily over 128-channel blocks. After quantising one block, GPTQ updates the remaining
672 full-precision weights with a rank- r approximation of the corresponding Hessian inverse, cheaply
673 computed from second-order activation statistics. This strategy yields near-optimal INT4 weights
674 with negligible calibration cost.
675676 A.3 SMOOTHQUANT
677678 SmoothQuant jointly scales activations and weights so that both can be represented with the same
679 uniform INT8 format. For each output channel, a learned scale $\sigma_k > 0$ migrates range from
680 activations to weights:
681

682
$$\tilde{\mathbf{Y}} = (\mathbf{A} \odot \sigma^{-1})(Q(\mathbf{W} \odot \sigma))^\top.$$

683

684 The scales $\{\sigma_k\}$ are obtained by minimising the worst-case per-channel quantisation error across the
685 calibration set, typically using a few thousand tokens.¹ Once trained, the scales are fused into \mathbf{W} and
686 the model runs on standard INT8 kernels without auxiliary tensors or runtime re-scaling.
687688 **Implementation Notes.** All three methods adopt per-channel affine quantisation. AWQ and GPTQ
689 target 4-bit weights and retain FP16 activations, whereas SmoothQuant yields a fully INT8 model. We
690 keep the original hyper-parameters recommended by the respective authors to ensure reproducibility
691 across codebases.
692700
701 ¹We use 2 048 tokens, following Xiao et al. (2023).

702 **B PROMPT**
703704 **System Prompt**
705706 Your task is to solving mathematical problems by following these rules:
707

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.

720 **Problem Statement (from PRM-800K Dataset Training Subsets)**
721722 The product of the proper positive integer factors of n can be written as $n^{(ax+b)/c}$, where x is the
723 number of positive divisors n has, c is a positive integer, and the greatest common factor of the
724 three integers a , b , and c is 1. What is $a + b + c$? [Gold Answer: 1]
725726 **Response Sample(from PRM-800K Dataset Training Subsets)**
727

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 $n^{x/2}$, 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 $n^{x/2}/n$, which is $n^{(x-2)/2}$.
6. This means that the product of the proper factors of n can be written as $n^{(ax+b)/c}$, 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]

739 **Evaluation Prompt**
740741 Role = 'HUMAN',
742743 Prompt = "Your task is to solving mathematical problems by following these rules:
744

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.

755 Problem: problem"

756
757**Assessment System Prompt**758
759**You are a helpful assistant that identifies and classifies errors in mathematical reasoning steps.**

760

You will be given:

761

- **Problem Statement:** A math problem statement.
- **Answers:** The right answer and answer from full-precision model and quantized model. Which model's answer is correct.
- **Full-Precision Reasoning:** The reasoning steps and final answer from a full-precision model.
- **Quantized-Model Reasoning:** The reasoning steps and final answer from a quantized model.
- **Error Type Definition:** The definition and explanation of error types.

769

Your task:770
771

1. **Ground Truth Verification:** Compare both models' answers against the provided correct answer.

772

2. **Error Detection Protocol (Quantized Model):**

773

If the quantized model is incorrect:

774

1. Trace error origin using this hierarchy:

775

- Conceptual Errors: conceptual_misunderstanding, contextual_oversight
- Reasoning Errors: logical_reasoning_error
- Method Errors: procedural_error, formula_rule_error
- Execution Errors: computational_error, symbolic_manipulation_error

782

2. Identify first point of divergence from correct reasoning.

783

3. Classify using the most specific applicable type.

784

4. Provide step-specific evidence.

785

3. **Conflict Resolution:**

786

1. If multiple types apply, choose the earliest in the hierarchy.

787

2. If ambiguity persists, use procedural_error as default.

788

Return your analysis in the following JSON format strictly:

789

```
{
  "quantized_error_analysis": {
    "primary_error_type": [...],
    "error_step": 1,
    "explanation": "Short evidence from reasoning steps",
    "confidence_score": 0.7 // between 0.7 and 1.0
  }
}
```

790

791

C HUMAN ANNOTATION GUIDEBOOK

800

PURPOSE

801

This guideline specifies the manual verification protocol applied to *disagreement cases* that survive the automated evaluation pipeline—namely the expert-LLM judges and the five-model majority vote. Annotators produce the *final ground-truth verdict* (error type, error step, explanation, confidence) for every instance in which

802

803

804

805

806

807

808

809

- the majority vote conflicts with the baseline judge **DeepSeek-R1**, or
- a “passed” case is randomly drawn for audit ($\approx 2\%$ of all cases).

810

MATERIALS PROVIDED

811

812

813

814

815

816

817

818

819

820

ERROR-TYPE TAXONOMY

821

822

823

824

825

826

827

828

829

830

831

1. **Conceptual Errors:** conceptual_misunderstanding, contextual_oversight
2. **Reasoning Errors:** logical_reasoning_error
3. **Method Errors:** procedural_error, formula_rule_error
4. **Execution Errors:** computational_error, symbolic_manipulation_error

Earliest-precedence rule: when multiple labels apply, choose the first that appears in the above list.

832

ANNOTATION PROCEDURE

833

834

835

836

837

838

839

840

841

842

843

844

845

846

847

1. **Answer verification.** Confirm which model(s) yield the correct final answer. If both are wrong, mark the case `dual_failure`.
2. **Locate first divergence.** Read `fp_trace` and `gt_trace` in parallel and find the earliest step where the quantized trace deviates from valid reasoning.
3. **Review automated evidence.** Inspect the five judge outputs and majority-vote result.
4. **Decision.**
 - 4.1. Adopt the ensemble consensus if at least three judges agree *unless* compelling counter-evidence exists.
 - 4.2. Otherwise, perform an independent assessment using the taxonomy in Sec. C.3.
5. **Label assignment.** Record `primary_error_type`, `error_step` (1-indexed), `explanation` (≤ 40 words, quote the critical step), `confidence_score` (Sec. C.5).
6. **Quality flag.** Set `needs_second_opinion` = `true` if residual uncertainty remains.

848

CONFIDENCE-SCORE HEURISTIC

849

850

851

852

853

854

OUTPUT SCHEMA

855

856

Annotators create `human_verdict.json` using

857

858

859

860

861

862

863

```
{
  "quantized_error_analysis": {
    "primary_error_type": "procedural_error",
    "error_step": 4,
    "explanation": "Applied quadratic formula with sign error at step 4.",
    "confidence_score": 0.83
  }
}
```

864
865

DIMENSION DEFINITION

866

- **Conceptual Errors** occur when the model exhibits a fundamental misunderstanding of the underlying principles or relevant context of the problem. This can manifest either as a conceptual misunderstanding, where the core ideas or foundational theories are not correctly grasped, resulting in an erroneous approach or framing of the problem; or as contextual oversight, in which crucial situational constraints or domain-specific factors (such as physical boundaries or geometric limitations) are overlooked, significantly distorting the solution process and its outcome.
- **Method Errors** refer to inaccuracies stemming from the improper selection or application of mathematical methods or established procedural approaches. Specifically, procedural errors happen when prescribed sequences or standard algorithms are incorrectly executed or entirely skipped, causing incomplete or invalid solutions. Formula rule errors are another subtype, characterized by the misuse or misapplication of relevant mathematical theorems, formulae, or rules—such as applying a formula in an inappropriate context—which fundamentally undermines the validity of the resulting calculations or conclusions.
- **Execution Errors** arise during the process of mathematical computation and symbolic manipulation. They encompass computational errors involving incorrect arithmetic or algebraic operations, such as flawed summations, erroneous expansions, or factorization mistakes, thus jeopardizing the accuracy of final answers. Additionally, symbolic manipulation errors include improper handling or representation of symbolic expressions, variables, or transformations. This could involve mislabeling variables or misinterpreting symbolic forms, leading to an incorrect representation and subsequent solution of the problem.
- **Reasoning Errors** involve flaws in the logical flow of problem-solving. Specifically, logical reasoning errors occur when there is a breakdown in the reasoning process itself, such that inference steps either do not logically follow one another or omit essential connections. This causes a logical gap or disconnect between the initial premises and the eventual conclusion, rendering the derived solution fundamentally flawed or unsupported.

872

873

874

875

876

877

878

879

880

881

882

883

884

885

886

887

888

889

890

891

892

893

894

895

896

897

898

899

900

DECISION AIDS

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

- *Conceptual misunderstanding*: misstates theorem before algebra begins.
- *Contextual oversight*: ignores domain restrictions or boundary conditions.
- *Logical reasoning error*: unsupported logical jump.
- *Procedural error*: applies an inappropriate solution method.
- *Formula rule error*: violates algebraic/derivative rule.
- *Computational error*: arithmetic slip.
- *Symbolic manipulation error*: incorrect simplification of an expression.

QUALITY CONTROL & ETHICS

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

- Two Annotators work independently; no discussion of live cases.
- Evaluate reasoning quality, not model identity; avoid bias.
- Flag any toxic or sensitive content present in traces.
- Project leads re-annotate 2 % of “agree” cases and all needs_second_opinion cases; overall agreement < 95% triggers targeted review.

Note.—DeepSeek-R1 is designated the baseline judge owing to its highest pilot agreement with human experts.

918 **D CASE STUDY**
919920 **D.1 EXAMPLE OF EXECUTION ERRORS**
921922 In Case #93, the task was to find the value of c for which the circle defined by $x^2 - 10x + y^2 +$
923 $6y + c = 0$ has a radius of 1. The correct approach involves completing the square, resulting in
924 $(x - 5)^2 + (y + 3)^2 = -c + 34$, and solving $-c + 34 = 1$ to find $c = 33$. However, the GPTQModel
925 made a computational error by incorrectly rearranging the equation as $34 - c = 1$ leading to the wrong
926 answer $c = 35$. This error stemmed from mishandling the sign of c during algebraic manipulation,
927 highlighting the importance of careful sign management in equation solving. The case #93 shows the
928 **Execution Errors**.929 **Problem #93**
930931 For what value of c will the circle with equation $x^2 - 10x + y^2 + 6y + c = 0$ have a radius
932 of length 1?
933934 **[Gold Answer: 33]**935 **Response of LLaMA-3.1-8B-Instruct-BF16 (Vanilla)**936
937 1. To find the radius of a circle, I need to complete the square for both x and y terms in
938 the equation.
939 2. I can rewrite the equation as $(x^2 - 10x) + (y^2 + 6y) = -c$.
940 3. To complete the square for $x^2 - 10x$, I add $(10/2)^2 = 25$ to both sides.
941 4. To complete the square for $y^2 + 6y$, I add $(6/2)^2 = 9$ to both sides.
942 5. This gives $(x^2 - 10x + 25) + (y^2 + 6y + 9) = -c + 25 + 9$.
943 6. I factor the left side as $(x - 5)^2 + (y + 3)^2$.
944 7. The right side simplifies to $-c + 34$.
945 8. Therefore the equation becomes $(x - 5)^2 + (y + 3)^2 = -c + 34$.
946 9. The radius of the circle is the square root of the right-hand side.
947 10. Setting the radius to 1 gives $-c + 34 = 1^2 = 1$.
948 11. Solving for c yields $c = 34 - 1 = 33$.949 **[Answer: 33]**950 **Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQ Model)**951
952 1. To find the radius, I first complete the square for the x and y terms and rewrite the
953 equation in the form $(x - h)^2 + (y - k)^2 = r^2$, where (h, k) is the centre and r is
954 the radius.
955 2. I begin with the x terms, adding and subtracting the same value to form a perfect
956 square.
957 3. Because the coefficient of x is -10 , I add and subtract $(-10/2)^2 = 25$ to obtain
958 $x^2 - 10x + 25 - 25$.
959 4. Factoring gives $(x - 5)^2$ and the equation becomes $(x - 5)^2 - 25 + y^2 + 6y + c = 0$.
960 5. I repeat the process for the y terms, adding and subtracting $(-6/2)^2 = 9$ to get
961 $y^2 + 6y + 9 - 9$.
962 6. Factoring yields $(y + 3)^2$ and the equation is now $(x - 5)^2 - 25 + (y + 3)^2 - 9 + c = 0$.
963 7. Combining constants, I have $(x - 5)^2 + (y + 3)^2 - 34 + c = 0$.
964 8. Setting $r^2 = 1$ requires $-34 + c = 1$; solving gives $c = 35$.
965 9. Hence $c = 35$.966 **[Answer: 35]**

972 D.2 EXAMPLE OF REASONING ERRORS
973

974 In Problem #128, the GPTQModel made a logical reasoning error when calculating the minimum
975 number of miles Suzanne could walk in February. While the model correctly identified that February
976 has 28 or 29 days, it missed the 27th day when calculating her walking schedule every third day,
977 leading to an incorrect conclusion of 8 walking days and 32 miles. In reality, the correct number of
978 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
979 highlights the model's logical reasoning gap in iterating through sequential intervals accurately. The
980 case #128 shows the **Reasoning Errors**.

981 Problem #128
982

983 **Suzanne walks four miles every third day. What is the fewest number of miles she can
984 walk in February?**

985 **Gold Answer: 36**

987 Response of LLaMA-3.1-8B-Instruct-BF16 (Vanilla)
988

- 989 To find the fewest number of miles Suzanne can walk in February, I need to figure
990 out how many days she walks four miles in that month.
- 991 February has 28 or 29 days, depending on whether it is a leap year or not.
- 992 I do not know if this is a leap year, so I will assume it has 28 days.
- 993 If Suzanne walks four miles every third day, she walks on the 3rd, 6th, 9th, 12th,
994 15th, 18th, 21st, 24th, and 27th of February.
- 995 That is nine days, so she walks a total of $9 \times 4 = 36$ miles in February.
- 996 This is the fewest number of miles she can walk; if she skipped any of those days
997 she would have to walk on a different day, increasing her total.
- 998 7. Therefore, the answer is 36 miles.

999 [Answer: 36]

1000 Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQ Model)
1001

- 1002 To find the fewest number of miles Suzanne can walk in February, I must determine
1003 how many days she walks during that month.
- 1004 February has 28 days, but some years have 29 days.
- 1005 I will assume February has 28 days, since that is the most common case.
- 1006 If Suzanne walks every third day, she will walk on days 3, 6, 9, 12, 15, 18, 21, and
1007 24.
- 1008 That is a total of 8 days, and she will walk 4 miles on each of those days.
- 1009 6. Thus, the fewest number of miles she can walk in February is $8 \times 4 = 32$ miles.

1010 [Answer: 32]

1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025

1026 D.3 EXAMPLE OF NO ERROR
10271028 Despite the care taken in designing our extraction scripts, a small number of predictions remain hard
1029 to classify, leading to a residual "**No Errors**" category. A strong judge model can usually flag these
1030 edge cases; for instance, in case #3812, Deepseek-R1 correctly returns **No Errors** after a meticulous
1031 comparison. We subsequently review such instances and update the final labels accordingly.1032 Most ambiguities stem from multiple notations for the same numeric value—particularly decimals
1033 versus fractions and natural versus programming syntax—e.g., $1/2$, 0.5 , $\frac{1}{2}$, $5E-01$,
1034 and 5×10^{-1} .
10351036 Problem #3812
10371038 **What is the probability that when a fair 6-sided die is rolled, a prime number faces up?**
10391040 **Gold Answer:** $\frac{1}{2}$ 1041
1042 Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQ Model)
10431044 1. To find the probability of rolling a prime number on a fair 6-sided die, I need to
1045 know how many prime numbers appear on such a die and how many faces there are
1046 in total.
1047 2. A prime number is a positive integer divisible only by 1 and itself. The prime
1048 numbers on a 6-sided die are 2, 3, and 5.
1049 3. There are 3 prime numbers on a 6-sided die, so the probability of rolling a prime
1050 number is 3 out of 6, or $1/2$.
1051 4. Converting this fraction to a decimal, I divide 1 by 2 to obtain 0.5.
1052 5. Alternatively, multiplying 1 by 100 and dividing by 2 gives 50%.
1053 6. Therefore, the probability of rolling a prime number on a fair 6-sided die is 0.5 or
1054 50%.
1055 [Answer: 0.5]

1056

1057

1058 Judgement of Deepseek-R1 (Primary Expert Model)
10591060 The quantized model's answer (0.5) is numerically equivalent to the correct answer (1/2). All
1061 reasoning steps are logically valid and mathematically correct, with no conceptual, procedural,
1062 or computational errors detected. [No_Error: []]

1063

1064

1065 D.4 EXAMPLE OF CONFLICTING JUDGMENTS
10661067 Case #342 illustrates our conflict-resolution protocol when the judge models disagree on an error
1068 label. We perform a five-way cross-model validation using Deepseek-R1, Deepseek-V3, ChatGPT,
1069 GPT-4o, and Qwen-Max. Each model independently assigns an error category to the quantized trace,
1070 and the final label is set by majority vote.1071 If the vote is inconclusive (e.g., a 2-2-1 split), a human annotator re-examines the example. The
1072 annotator consults (i) each model's confidence score, (ii) the accompanying explanations, and (iii)
1073 the step-by-step reasoning provided by Deepseek-R1, together with the raw model outputs. This
1074 double-check ensures that every ambiguous case receives a consistent, well-justified error type.
1075

1076

1077

1078 Problem #342
10791078 **What integer x satisfies $\frac{1}{4} < \frac{x}{7} < \frac{1}{3}$?**1079 **Gold Answer:** 2

1080
1081

Response of LLaMA-3.1-8B-Instruct-W4A16 (GPTQ Model)

1082
1083
1084
1085

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

[Answer:

2,3,4,5,6]

1100

1101
1102
1103

Judgement of Deepseek-R1 (Primary Expert Model)

1104
1105
1106
1107

When multiplying $4/3$ by 7, the correct result is $28/3 \approx 9.33$. In step 5 the quantized model treats $7 \times (4/3)$ as $28/3$ but, in step 7, fails to divide this by 4 correctly, giving an upper bound of 7 instead of $7/3 \approx 2.33$. [Execution_Errors: computational_error]

1108
1109
1110
1111
1112
1113

Judgement of Deepseek-V3

The quantized model incorrectly multiplies both sides by 4 instead of 7, the proper common denominator. This leads to a procedural error that propagates through the subsequent steps. [Method_Errors: procedural_error]

1114
1115

Judgement of ChatGPT

1116
1117
1118
1119

The model simplifies the inequality to $\frac{7}{4} < x < \frac{7}{1}$, overlooking that the upper limit should be $(\frac{28}{3}) \div 4 = \frac{7}{3}$, not $\frac{7}{1}$. [Conceptual_Errors: contextual_oversight]

1120
1121
1122
1123
1124

Judgement of GPT-4O

The quantized model reports the range $\frac{7}{4} < x < \frac{7}{1}$ instead of the correct $\frac{7}{4} < x < \frac{7}{3}$, yielding an incorrect set of integer solutions. [Method_Errors: procedural_error]

1125
1126

Judgement of Qwen-Max

1127
1128
1129
1130
1131
1132
1133

The model simplifies $\frac{7}{4} < x < \frac{28}{3}$ to $1.75 < x < 7$. The upper bound $\frac{28}{3}$ is approximately 9.33, not 7. [Conceptual_Errors: conceptual_misunderstanding]

1134

E EXPERIMENT RESULTS

1135

E.1 QUANTIZATION ON LARGER MODELS

1136 To strengthen the generality of our conclusions and bringing potential insights, we have conducted
 1137 additional experiments on the Qwen3 series, which natively support a "thinking mode" as the
 1138 default configuration. These models already integrate internal CoT-like mechanisms, making them
 1139 representative of both standard and "thinking" (CoT) LLMs. We evaluated Qwen3-8B, Qwen3-14B,
 1140 and Qwen3-32B under identical quantization settings (bit-width, calibration set, and hyperparameters)
 1141 as other quantization methods.
 1142

1143 Table 3: Performance of Qwen3 models on MATH under different quantization methods.

	Qwen3-8B			Qwen3-14B			Qwen3-32B		
	Van.	AWQ	GPTQ	Van.	AWQ	GPTQ	Van.	AWQ	GPTQ
MATH	55.88	53.96	53.58	63.52	62.28	63.12	66.92	65.72	65.94

1144 *Note: Van. denotes the vanilla full-precision baseline.*

1145 These results yield several key insights:

- 1146 • **Larger models show greater quantization robustness.** Accuracy degradation from full-
 1147 precision to quantized versions diminishes significantly as model size increases, consistent
 1148 with observations from other studies. We believe because larger models possess a richer
 1149 parameter space, which grants stronger robustness to quantization-induced numerical errors
 1150 when mapping from high-precision to low-precision data formats.
- 1151 • **Mild regularization effects appear in quantized models.** You can find an interesting
 1152 result that Qwen3-14B under GPTQ slightly outperforms its vanilla counterpart, suggesting
 1153 that moderate compression may enhance generalization, a phenomenon also noted in recent
 1154 quantization studies. Notably, this effect has also been frequently cited as one of the
 1155 underlying reasons why Quantization-Aware Training (QAT) can sometimes improve the
 1156 generalization ability of post-quantized models.
- 1157 • **Scalability of the "Silver Bullet" principle.** The consistently small degradation across
 1158 larger models supports our hypothesis that targeted recovery using compact, well-curated
 1159 data is even more effective for high-capacity models with stronger learning abilities.

1160

E.2 ADDITIONAL EVALUATION ON GENERAL-REASONING BENCHMARKS

1161 To complement our analysis on mathematical reasoning, we additionally evaluate Qwen2.5-
 1162 0.5B/1.5B/3B/7B-Instruct models on several widely used benchmarks that probe general science
 1163 question answering, commonsense reasoning, and instruction following:

- 1164 • **ARC Easy and ARC Challenge:** a science exam multiple-choice benchmark (grades 3–9)
 1165 with two difficulty splits. Most questions provide four options and the Challenge split
 1166 requires more complex reasoning (Clark et al., 2018).
- 1167 • **HellaSwag:** a commonsense inference benchmark with 70k multiple-choice questions. Each
 1168 item provides a scenario and four possible continuations; the distractors are adversarially
 1169 generated to fool models while remaining trivial for humans (Zellers et al., 2019).
- 1170 • **IFEval:** an instruction-following benchmark built from verifiable constraints (for example
 1171 “write more than 400 words”) that focuses on controllable instructions and reduces the bias
 1172 of LLM-based judges (Zhou et al., 2023).
- 1173 • **CommonSenseQA:** a 12k-question multiple-choice benchmark that requires various forms
 1174 of commonsense knowledge, with one correct answer and four distractors per question (Tal-
 1175 mor et al., 2019).

1176 For each benchmark we compare the vanilla full-precision model with its AWQ and GPTQ quantized
 1177 counterparts, and summarize the average degradation in Table 4. Compared with the larger drops
 1178 on mathematical reasoning tasks, the performance drops on general commonsense and language-
 1179 understanding benchmarks are substantially smaller, both in absolute score reduction and in relative
 1180

1188 Table 4: Average accuracy degradation of Qwen2.5-Instruct models on additional benchmarks
 1189 under post-training quantization. Negative values indicate performance drop compared with the
 1190 corresponding full-precision models.

Benchmark	Type	Avg. accuracy drop ↓ (points)	Avg. relative drop ↓ (%)
ARC-c	General science QA	-3.56	-4.10
ARC-e	General science QA	-3.59	-4.76
CommonSenseQA	Commonsense QA	-3.38	-5.06
HellaSwag	Commonsense	-2.35	-3.95
IFEval	Instruction following	-2.61	-5.01
GSM8K	Grade-school math	-6.78	-13.21
MATH	Competition math	-10.56	-29.84

1201 Table 5: Accuracy (%) of Qwen2.5-Instruct models on general benchmarks before and after quantiza-
 1202 tion.

Model Scale	Method	ARC-c	ARC-e	CommonSenseQA	HellaSwag	IFEval
0.5B	Vanilla	46.44	65.43	59.38	39.43	36.69
	AWQ	51.86	64.90	53.81	37.81	36.81
	GPTQ	43.39	50.97	50.37	36.47	32.01
1.5B	Vanilla	77.97	89.95	76.00	62.19	50.96
	AWQ	72.54	84.13	73.38	58.62	47.60
	GPTQ	71.86	86.42	71.42	60.19	46.04
3B	Vanilla	84.75	91.53	78.62	76.62	68.23
	AWQ	75.59	88.36	76.41	73.28	65.47
	GPTQ	78.64	89.77	77.56	73.52	64.87
7B	Vanilla	86.10	92.59	84.19	85.18	77.82
	AWQ	86.44	93.47	82.47	84.38	76.86
	GPTQ	84.75	92.24	83.95	83.75	76.86

1217 *Note: Inst. is an abbreviation of Instruct.*

1218 percentage. This supports our claim that post-training quantization disproportionately affects mathe-
 1219 matical reasoning ability while having only mild impact on general language capabilities.

1220 The detailed per-model results are listed in Table 5. We report accuracy for each Qwen2.5-Instruct
 1221 checkpoint and for each quantization method.

1224 E.3 CASE STATISTICS

1225 Table 6 shows detailed statistics of all error types. The total number of cases varies slightly across
 1226 models due to differences in error rates and scores.

1228 E.4 SUBJECT-WISE AND DIFFICULTY-WISE DEGRADATION ON MATH

1229 To better understand how quantization-induced degradation relates to problem structure, we leverage
 1230 the rich annotations in the MATH dataset, which covers multiple subject domains (such as algebra,
 1231 geometry, number theory and combinatorics) and five difficulty levels (Level 1 to Level 5). For each
 1232 model scale and quantization method, we compute the distribution of errors across mathematical
 1233 subfields and difficulty levels. This analysis connects quantization sensitivity with different types of
 1234 reasoning and provides practical guidance for deploying quantized models under varying reasoning
 1235 complexities.

1236 From the subject-wise analysis in Table 7, we observe a clear and consistent trend across all model
 1237 scales: quantization-induced degradation is not uniformly distributed over mathematical domains.
 1238 Subfields that involve multi-step symbolic manipulation, such as Intermediate Algebra, Precalculus
 1239 and more advanced algebraic transformations, show noticeably larger performance drops for both
 1240 AWQ and GPTQ. In contrast, domains that rely more on direct recall or simpler numerical reasoning
 1241 remain comparatively stable. This pattern suggests that low-bit perturbations disproportionately affect

Table 6: Detailed statistics of all error types. The total number of cases varies slightly across models due to differences in error rates and scores. For full-precision models, all incorrectly answered problems are included; for quantized models, only those problems solved correctly by the full-precision model but failed after quantization are counted.

	Method	Conceptual Errors	Method Errors	Reasoning Errors	Execution Errors	No Error	TTL
Llama-3.1-8B-Inst.	Van.	1622	313	427	380	28	2770
	AWQ	286	86	5	136	4	517
	GPTQ	310	97	5	128	0	540
	SQ	199	58	7	102	3	369
Llama-3.2-3B-Inst.	Van.	1760	369	387	387	82	2985
	AWQ	317	91	1	107	1	517
	GPTQ	326	102	5	123	2	558
	SQ	236	65	4	88	3	396
Llama-3.2-1B-Inst.	Van.	2521	515	491	491	40	4058
	AWQ	287	87	6	108	0	488
	GPTQ	315	104	2	85	1	507
	SQ	196	85	4	70	0	355
Qwen2.5-7B-Inst.	Van.	872	324	290	303	44	1833
	AWQ	262	72	13	103	1	451
	GPTQ	267	82	11	116	4	480
	SQ	183	53	5	42	9	292
Qwen2.5-3B-Inst.	Van.	1217	322	299	362	40	2240
	AWQ	386	93	7	139	2	627
	GPTQ	351	120	7	130	3	611
	SQ	225	65	11	84	2	387
Qwen2.5-1.5B-Inst.	Van.	1937	273	445	373	49	3077
	AWQ	344	76	8	93	0	521
	GPTQ	344	82	2	106	1	535
	SQ	185	53	2	56	0	296
Qwen2.5-0.5B-Inst.	Van.	2834	406	264	312	104	3920
	AWQ	429	89	4	96	1	619
	GPTQ	521	59	3	70	1	654
	SQ	183	53	5	42	9	292

Notes: Van. denotes the vanilla full-precision baseline; Inst. is an abbreviation of Instruct.

Table 7: Distribution of errors across mathematical domains on MATH under different quantization settings. All values are percentages. All experiments are conducted on Qwen2.5-Instruct models.

Model	Method	Number Theory	Counting & Prob.	Interm. Algebra	Algebra	Geometry	Prealgebra	Precalculus
0.5B	Van.	11.59	10.33	20.48	18.65	10.30	16.02	12.63
	AWQ	11.26	9.82	18.77	22.24	9.84	16.68	11.39
	GPTQ	11.08	9.80	18.74	22.57	9.65	16.71	11.45
1.5B	Van.	11.31	10.67	22.98	15.79	10.67	13.73	14.86
	AWQ	11.39	10.26	20.16	19.16	10.62	15.72	12.68
	GPTQ	11.57	10.28	21.74	18.99	10.28	13.94	13.56
3B	Van.	10.40	10.35	24.85	12.91	11.94	12.07	17.49
	AWQ	10.74	10.78	25.05	13.92	11.13	12.29	16.09
	GPTQ	11.23	9.98	24.49	14.58	11.20	12.30	16.22
7B	Van.	7.96	10.31	26.28	11.75	12.87	11.91	18.91
	AWQ	9.58	9.88	25.89	11.87	12.22	11.67	18.90
	GPTQ	8.56	10.50	27.13	11.20	12.39	11.60	18.62

Note: Van. denotes the vanilla full-precision baseline.

tasks that require long-range dependency tracking and precise arithmetic transformations, which is consistent with the step-level error analysis in the main paper.

From the difficulty-level analysis in Table 8, we see a monotonic increase in degradation as problem difficulty grows. Level 1 and Level 2 questions exhibit only marginal changes after quantization, whereas degradation becomes much more pronounced for Levels 3 to 5. For Level 5 in particular, the gap between full-precision and quantized models can exceed 6–10 percentage points even for larger models.

These findings provide empirical evidence for the following points:

1296 Table 8: Distribution of errors across difficulty levels on MATH under different quantization settings.
 1297 All values are percentages.
 1298

Model	Level 1	Level 2	Level 3	Level 4	Level 5
Qwen2.5-0.5B-Inst. Van.	5.72	15.41	21.39	26.04	31.44
Qwen2.5-0.5B-Inst. AWQ	7.53	17.29	22.43	24.99	27.75
Qwen2.5-0.5B-Inst. GPTQ	7.17	17.35	22.44	24.99	28.05
Qwen2.5-1.5B-Inst. Van.	4.38	10.42	20.30	25.59	35.71
Qwen2.5-1.5B-Inst. AWQ	4.89	16.57	21.37	25.31	31.48
Qwen2.5-1.5B-Inst. GPTQ	4.41	14.22	21.03	23.82	36.53
Qwen2.5-3B-Inst. Van.	3.39	12.20	19.03	25.73	39.65
Qwen2.5-3B-Inst. AWQ	3.30	12.21	20.12	25.86	38.50
Qwen2.5-3B-Inst. GPTQ	3.16	15.27	19.52	26.50	35.55
Qwen2.5-7B-Inst. Van.	3.42	12.45	19.07	26.01	39.05
Qwen2.5-7B-Inst. AWQ	3.34	12.02	18.75	25.79	40.10
Qwen2.5-7B-Inst. GPTQ	3.38	12.00	19.31	25.19	40.12

1312 *Notes: Van. denotes the vanilla full-precision baseline; Inst. is an abbreviation of Instruct.*

1313

- 1314 • Conceptually demanding and algebraically heavy subfields of MATH are especially vulnerable to precision reduction.
- 1315 • Higher-difficulty problems, which require deeper chains of reasoning, tend to accumulate quantization noise and errors more severely.

1316
 1317 Overall, this analysis further supports our claim that post-training quantization disproportionately
 1318 affects mathematical reasoning compared with general language understanding. We include these
 1319 tables and observations in the appendix for completeness.

1320
 1321
 1322
 1323
 1324
 1325
 1326
 1327
 1328
 1329
 1330
 1331
 1332
 1333
 1334
 1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349

1350 F CAPABILITY RESTORATION RESULTS

1352 Table 9: Capability restoration results on GSM8K, MATH500, and MMLU benchmarks across
 1353 different model scales using our curated *Silver Bullet* datasets. **Full Precision** refers to the full-
 1354 precision model after format alignment. **BF** indicates performance before restoration, while **AF**
 1355 shows performance after applying our restoration pipeline.

1357 Quantization	1358 Task	1359 Llama-3-Inst.			1360 Qwen2.5-Inst.		
		1361 1B	1362 3B	1363 8B	1364 0.5B	1365 1.5B	1366 3B
1358 Full Precision	GSM8K	38.44	71.34	76.88	42.99	61.87	76.04
	MATH500	18	32.4	36.4	16.6	22.2	39
	MMLU	45.14	61.81	68.62	45.49	59.71	65.1
	AVG	33.86	55.18	60.63	35.03	47.93	60.05
1361 AWQ-BF	GSM8K	35.03	70.58	77.1	27.9	53.15	70.36
	MATH500	13.8	29.6	33.2	8.2	21	29
	MMLU	43.26	60.08	67	42.65	57.65	63.16
	AVG	30.70	53.42	59.10	26.25	43.93	54.17
1362 GPTQ-BF	GSM8K	32.15	69.67	76.27	25.02	57.09	68.54
	MATH500	15.4	26.4	33.6	8.6	21.8	31.2
	MMLU	42.07	59.49	66.44	42.91	57.86	62.09
	AVG	29.87	51.85	54.94	25.51	45.58	53.94
1363 AWQ-AF	GSM8K	40.49	74.3	80.14	26.38	56.86	68.84
	MATH500	15.2	36.8	34.6	9.4	26.4	38.4
	MMLU	43.72	60.57	67.67	43.99	59.43	64.8
	AVG	33.14	57.22	60.80	26.59	47.56	57.35
1364 GPTQ-AF	GSM8K	37.83	73.01	79.68	25.93	55.65	75.21
	MATH500	18.2	33	36	8.4	25.2	40.6
	MMLU	42.29	59.9	67.23	44.15	59.43	63.63
	AVG	32.77	55.30	60.97	26.16	46.76	59.81

1372 As shown in Table 9, after capability restoration using our *Silver Bullet* dataset, the quantized 4-
 1373 bit models not only recover but even surpass the performance of their full-precision counterparts
 1374 on the MATH benchmark. Meanwhile, performance on GSM8K remains stable, and accuracy
 1375 on MMLU—a diverse benchmark covering various complex reasoning tasks—is also preserved.
 1376 These results demonstrate that our *Silver Bullet* dataset effectively restores mathematical reasoning
 1377 capabilities without compromising general-purpose abilities, highlighting both the effectiveness and
 1378 generalizability of our approach.

1380 G THE USAGE OF LLM

1382 In this work, Large Language Models (LLMs) were used as auxiliary tools to support our research
 1383 process, but not to generate novel scientific content. Specifically, their usage includes:

- 1385 • **Editing and polishing.** LLMs were employed for minor grammar checking, improving
 1386 clarity, and rephrasing sentences for readability in the manuscript. All scientific ideas,
 1387 methodology, and experiments were designed and written by the authors.
- 1388 • **Facilitating annotation.** During the construction of our automated error-assessment
 1389 pipeline, LLMs were used as expert judges to classify error types in reasoning traces.
 1390 Their outputs were combined via majority voting and, when necessary, verified by human
 1391 annotators to ensure reliability.
- 1392 • **Experiment assistance.** LLMs were queried to simulate baseline reasoning traces for
 1393 building our contrastive “Silver Bullet” datasets, which were later curated, filtered, and
 1394 validated by the authors. This step complements human effort by accelerating the generation
 1395 of positive examples.

1396 We emphasize that all key contributions—including research ideas, methodology design, experimental
 1397 execution, and analysis—were conceived and implemented by the authors.