AQLoRA: A Zero-Search Recipe for Fast Quantized LoRA Fine-Tuning

Md Romyull Islam *mislam22@students.kennesaw.edu*
Kennesaw State University

arXiv:2608.23816v2 [cs.LG] 11 Oct 2026

Abstract

Quantized fine-tuning (QLoRA) saves memory but not time. It dequantizes every 4-bit weight on the fly, so it trains more slowly than fp16 LoRA. We present AQLoRA (Adaptive-Quantization LoRA), a recipe that buys part of that time back. One CPU pass over the weights fixes both the precision map and the adapter placement, with no search and no calibration data. The pass ranks layers by NF4 reconstruction error and keeps the top-K in fp16, where K is a fixed fraction of the linear layers. Those layers skip dequantization, which is where the speed comes from. A quality setting adapts every layer. A speed setting adapts only the top blocks, so the backward pass stops early. The rule matches the accuracy of Unsloth's hand-curated dynamic-4bit selection, in seconds, where search-based allocation needs repeated passes over calibration data.

We evaluate on Commonsense-170K across six models and four architecture families, from 1.4B to 14B. The speed setting trains $11.3 \pm 2.3\%$ faster than well-tuned QLoRA and gives up 0.1 to 1.5 accuracy points. It was faster in every one of thirteen independent timing sessions, at worst by 7%. The quality setting trains $4.8 \pm 2.0\%$ faster. Its accuracy is level with QLoRA on every model and within a point of fp16 LoRA. The memory premium depends on how the architecture packs its layers: 0.06 to 0.8 GiB on the 1–4B models, 0.5 GiB at 7B, 1.9 GiB at 14B. These error bars are measured between independent sessions, not within one. Three rules follow for timing on shared hardware. Fix the measurement duration, not the step count. Measure the noise floor from a duplicated arm, not a nearly identical method. Repeat whole sessions: on one model, three sessions that each passed the floor test spread over 8.2 points, and the random seed controls almost none of it.

We validate the recipe with controls and report the two that failed. Choosing adapter layers by weight density is no better than random. Choosing protected layers by quantization error is not either. We test that on five models at three budgets, matched on protected memory, and downstream on two of them. The number of protected layers sets the speed; which layers they are does not measurably change accuracy.

Code and the raw per-run artifacts: `https://github.com/Romyull-Islam/AQLoRA`

1 Introduction

Low-rank adaptation (LoRA) (Hu et al., 2022) and its quantized form QLoRA (Dettmers et al., 2023) made large-language-model fine-tuning cheap on memory. A frozen 4-bit NF4 backbone with a small adapter fits a billion-parameter model on one commodity GPU. Memory savings are not speed savings. QLoRA dequantizes every 4-bit weight block back to half precision on each forward and backward pass, and that cost makes it slower than fp16 LoRA. Several studies put NF4 QLoRA at 0.7 to 0.8 times the training throughput of plain LoRA (Ko, 2025; Khaki et al., 2025; Huang et al., 2025). People who quantize to fit a model at all, whether on an older datacenter card, a laptop, or a phone, then pay a speed penalty on top of the memory saving.

Part of that penalty can be bought back. An fp16 layer in an otherwise-quantized model skips dequantization, so every layer kept in fp16 converts memory into speed. On phi-1.5, 0.22 GiB of fp16 layers makes training

6.3% faster than QLoRA (Table 4). The natural question is which layers to keep, and the sensitivity literature suggests an answer: NF4 error concentrates unevenly across layers, so keep the most damaged ones. We adopt that rule and test it. The count of protected layers drives the speed. Which layers are protected does not measurably change accuracy (§7.2). The rule survives on different grounds, being deterministic, data-free, and as accurate as a selection industry already ships by hand (89.0 against 89.0 on Qwen2.5-7B, §7.2). Layer protection itself is already in use. Unsloth's dynamic 4-bit models (Unsloth team, 2024) keep chosen layers in 16-bit to restore QLoRA accuracy, and QEFT (Lee et al., 2024a) keeps sensitive columns in fp16 with custom kernels and reports both quality and speed gains. What is missing is a stated rule, free of calibration data, that runs on the stock HuggingFace and bitsandbytes stack with no custom kernels. Also missing is a comparison against the search-based bit-and-rank allocators (Zhou et al., 2025a; 2026) that define the rest of this space.

AQLoRA fills that gap. One CPU pass over the weights takes 5 to 37 s on the models here, with no data and no forward pass. It produces two maps. A precision map keeps the top-K most NF4-damaged layers in fp16 within a layer budget. A placement map adapts every layer for quality, or only the top blocks for speed, so the backward pass stops early. It can be combined with rank-stabilized scaling (Kalajdzievski, 2023), though that combination is model-dependent (Appendix D). We apply every generic speed lever (length-grouped batching, a fused optimizer, the attention backend) to all baselines alike, so any gain we measure comes from AQLoRA's own parts and not from tuning.

On Commonsense-170K with the 8-task benchmark, across six models and four architecture families from 1.4B to 14B, AQLoRA's quality setting matches QLoRA everywhere and comes within a point of the fp16 LoRA reference. The same parity holds on GSM8K math reasoning. Accuracy is a parity result, not a win; the gain is speed. Sensitivity-based protection and early backward stopping have prior work, which we cite. Our part is the recipe that composes them, the controls on it, and a speed result that holds where kernel- and FP8-based methods cannot run. We also checked a weight-density rule for which layers to adapt, found it no better than random, and report that. A rank sweep ($r \in \{4, 8, 16, 32\}$) confirms the method ordering is not an artifact of one adapter size. It also turns up a failure worth flagging: rsLoRA, the strongest baseline at r=8, collapses to chance accuracy at r=32 under the shared learning rate while its training loss still looks healthy.

Contributions.

1. A **budgeted protection recipe** for quantized LoRA fine-tuning: one calibration-free CPU pass fixes which layers stay in fp16 and which get adapters, from a budget on the fraction of layers. Sensitivity-ranked protection is established in inference PTQ (Wu et al., 2020; Dong et al., 2020; Dumitru et al., 2024) and industry practice (Unsloth team, 2024). Ours is the formalization and its evaluation against Unsloth's deployed hand-picked selection on accuracy. Unlike search-based allocation (Zhou et al., 2025a; 2026; Guo et al., 2024), it needs no forward pass. The ranking inside it is a tie-break we test in item 2, not a contribution we claim.

2. **Two negative results from our own controls, and the speed mechanism that remains**: a weight-density rule for which layers to adapt is no better than random, and neither is our NF4-error ranking for which layers to protect. We test the ranking on five models at three budgets against random sets matched on protected parameters, and on downstream accuracy on two of them. The test also corrected our own control: random draws matched on layer count are not matched on memory where layer sizes vary, and that artifact had reversed the result on one model. What does help is adapting only the top blocks, the known backward-stopping speedup (Rücklé et al., 2021) on a 4-bit base; we measure it and do not claim it as new.

3. A **systems result on the stock stack**, on unmodified HuggingFace, bitsandbytes, and PEFT, timed over thirteen sessions on five models. The speed setting trains $11.3 \pm 2.3\%$ faster than well-tuned QLoRA for 0.04 to 0.8 GiB more peak memory, giving up 0.1 to 1.5 accuracy points. The quality setting trains $4.8 \pm 2.0\%$ faster and stays within a point of fp16 LoRA accuracy. Both are faster than QLoRA in every session. This holds on V100-class hardware where FP8 methods (Choi et al., 2025) do not run.

4. A **fair-lever evaluation protocol**: every universal efficiency lever (length-grouped batching, fused optimizer, dropout, attention backend) is applied to *every* baseline, so the measured advantage comes only from AQLoRA's own mechanisms. We release the code and the raw per-run artifacts, so every number in the paper can be recomputed without a GPU. The full lever ladder (Appendix L) makes each lever's contribution auditable.

2 Related Work

Quantized PEFT. QLoRA (Dettmers et al., 2023) backpropagates through a frozen NF4 base into LoRA (Hu et al., 2022). Successors improve the quantization-aware initialization (LoftQ (Li et al., 2024), ApiQ (Liao et al., 2024), CLoQ (Deng et al., 2025), QERA (Zhang et al., 2025a)), mergeability (QA-LoRA (Xu et al., 2024)), information retention (IR-QLoRA (Qin et al., 2024)), or quantizer modularity (ModuLoRA (Yin et al., 2024), TMLR). None of these targets training *speed*.

Mixed precision and sensitivity-based allocation. Ranking layers by quantization sensitivity and keeping the worst in higher precision is long-standing PTQ practice (Wu et al., 2020; Dong et al., 2020; Dumitru et al., 2024). It is deployed at scale by llama.cpp's k-quant mixes and by Unsloth's dynamic 4-bit models (Unsloth team, 2024), which keep chosen layers in 16-bit for QLoRA fine-tuning based on measured quantization error, though without a stated selection rule or a systematic study. QEFT (Lee et al., 2024a) and OWQ (Lee et al., 2024b) keep sensitive *columns* in fp16 with custom layouts and report both quality and fine-tuning speed gains. Another line assigns per-layer bits and ranks by search: LQ-LoRA uses an ILP (Guo et al., 2024), LowRA an ILP assigner with learned mappings that reaches 1.15 bits (Zhou et al., 2025c), QR-Adaptor a gradient-free search (Zhou et al., 2025a), and AutoQRA an evolutionary and Bayesian search (Zhou et al., 2026). AQLoRA differs from these in scope, not idea. It works at whole-layer granularity on the stock stack rather than QEFT's kernels, states a budgeted rule where Unsloth ships a hand-picked list, and profiles in 5 to 37 s where the others search.

Adapter placement and selective backpropagation. That not every layer needs an adapter dates to AdapterDrop (Rücklé et al., 2021), which removed lower-layer adapters and quantified the resulting backward-pass savings. LoRA-drop (Zhou et al., 2025b), SR-LoRA (Zhang et al., 2025b), PaFi (Liao et al., 2023), and Aletheia (Saket, 2026) select layers or parameters by output, spectral, magnitude, or probe statistics. AdaLoRA (Zhang et al., 2023) adapts rank budgets. Stochastic backward reduction (DropBP (Woo et al., 2024), LISA (Pan et al., 2024), LCSB (Park et al., 2026)), layer freezing (Wang et al., 2023; Huang et al., 2024), and side networks (Sung et al., 2022) all cut backward compute. AQLoRA takes depth-restricted placement from this line rather than inventing it, and adds the measured speed and quality trade-off on a 4-bit base, including how it interacts with gradient checkpointing.

Fast fine-tuning systems. Kernel-level systems (Unsloth, Liger, LoRAFusion (Zhu et al., 2026)) and FP8 methods (FALQON (Choi et al., 2025)) set the absolute speed bar on modern datacenter GPUs; SparseLoRA (Khaki et al., 2025) cuts compute via contextual sparsity. These need custom kernels or Hopper-class hardware. AQLoRA aims at the other case, the stock stack on pre-Hopper and edge devices, and its levers still compose with these systems.

On-device fine-tuning. Real-phone LLM fine-tuning exists via memory-efficient BP (MeBP (Song & Tang, 2025)), zeroth-order methods (ZeroQAT (Tan et al., 2025)), llama.cpp-based mobile LoRA (Tether AI, 2025), and engines like PockEngine (Zhu et al., 2023) and MobileFineTuner (Geng et al., 2025); EDGE-LLM (Yu et al., 2024) adapts compression in simulation. All fix base precision by a uniform setting or a fixed heuristic mix; none selects per-layer precision from a measurement of the model being tuned.

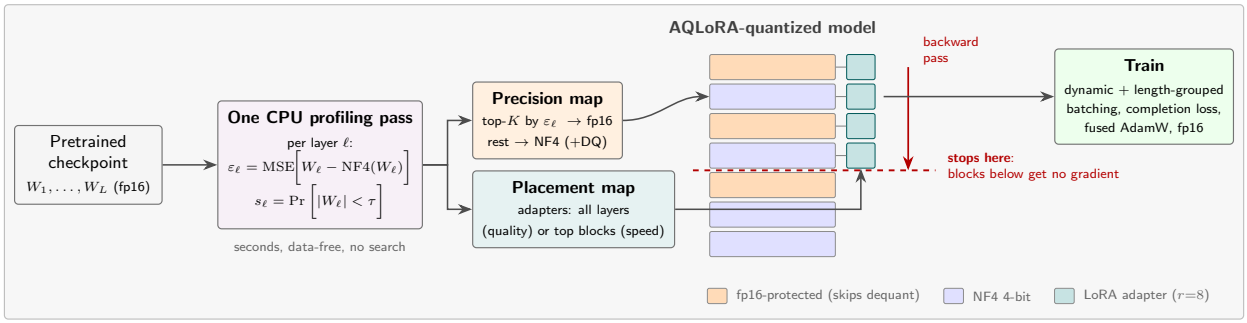


Figure 1: The AQLoRA pipeline.

3 Method

3.1 Profiling pass

For each linear layer ℓ with weights W_ℓ, flattened into blocks b of 64 elements (matching bitsandbytes),

$$\varepsilon_\ell \;=\; \frac{1}{|W_\ell|}\sum_{b}\left\|b - \text{absmax-NF4}^{-1}\big(\text{absmax-NF4}(b)\big)\right\|_2^2, \qquad s_\ell \;=\; \Pr\left[\,|W_{\ell,ij}| < \tau\,\right], \tag{1}$$

with τ=0.01 and the true 16-level NF4 quantile grid. The pass runs on CPU over the fp16 checkpoint, is data-free, and requires no forward pass. On 40 CPU threads with the checkpoint in the page cache, it takes 5 s on the 1B models, 12 to 13 s at 3 to 4B, 25 s at 7B and 37 s at 14B; reading the checkpoint from disk adds up to 21 s (released as `profile_timing.json`). This is negligible beside the multi-hour fine-tuning run. The calibration- and search-based allocation of QR-Adaptor and AutoQRA (Zhou et al., 2025a; 2026) instead evaluates each candidate configuration with forward passes on calibration data.

3.2 Precision allocation (zero-search)

The errors ε_ℓ are far from uniform. On phi-1.5 they span 11.8× across the 144 linear layers, and the top fifth hold 29% of the total (Figure 3, Appendix B), so a small fp16 budget can cover the worst offenders. Given a protection budget $\rho \in [0, 1]$ (`sensitive-frac`), the $K{=}\lceil \rho L \rceil$ layers with largest ε_ℓ are loaded in fp16; all others in NF4 with double quantization; `lm_head` stays fp16. Note that ρ counts layers, not parameters: the same ρ buys 4.3% of the weights on Llama-3.2-1B and 16.3% on Phi-3-mini (Table 8), so it is a budget on the rule's input and not directly on memory. The rule is a fixed top-K cut with no calibration data, no search, and no ILP. We call this *zero-search* in a specific sense: the configuration is fixed by one deterministic pass over the weights, with no evaluation of candidate configurations. It is not zero-cost. The pass itself costs the 5 to 37 s measured above, and the budget ρ remains a hyperparameter we set once and do not tune per model. It helps speed, since fp16 layers skip the per-step dequantization that is the main reason QLoRA runs slower than fp16 LoRA (Ko, 2025; Khaki et al., 2025).

We state one caveat here rather than leaving it to the ablations, because it concerns the rule itself. Ranking by ε_ℓ is the natural reading of the sensitivity literature, and it is what we implemented, but when we test it against parameter-matched random selections it is not measurably ahead (§7.2). The count of fp16 layers is what buys the speed, and on the evidence we have the identity of those layers is not what buys the quality. We keep the rule because it is deterministic, data-free, costs 5 to 37 s, and matches the accuracy of a deployed hand-curated selection, not because we can show it beats picking K layers arbitrarily.

3.3 Adapter placement

An adapter budget α (`adapter-frac`) sets how many of the L layers get an adapter. Adapting fewer layers is faster but costs accuracy (0.7 points at α=.7 on phi-1.5, §7.1), so α trades speed against accuracy. We considered choosing the $\lceil \alpha L \rceil$ layers by weight density (s_ℓ), but this does no better than a random

subset (§7.1), so we do not use it: the quality setting adapts every layer. For speed we instead use **depth** placement, which puts adapters on the top blocks only. No trainable parameter then sits below the cut block, so autograd stops the backward pass there, and with non-reentrant gradient checkpointing the lower blocks also skip recomputation. This is the AdapterDrop mechanism (Rücklé et al., 2021), and unlike a random subset it makes the saving concrete rather than incidental.

3.4 Composition and levers

AQLoRA can be combined with rank-stabilized scaling (Kalajdzievski, 2023) (+rs rows); Appendix D shows the combination is model-dependent. All experiments also use a set of shared levers, applied to every method: SDPA attention, LoRA dropout 0, length-grouped dynamic batching, and a fused AdamW. The lever ladder (Table 16) measures each one.

We report two settings of AQLoRA, both with ρ=.2 and no rsLoRA scaling (the rsLoRA composition is optional and model-dependent, §D). **AQLoRA-q** is the quality setting: adapters on every layer (α=1.0). **AQLoRA-s** is the speed setting: depth placement with α=.85 and early backward stopping. These are two settings of one budget, not separate methods.

4 Experimental setup

Protocol. Commonsense-170K fine-tuning with the Alpaca template and completion-only loss, evaluated by greedy generation on the eight held-out test sets (BoolQ, PIQA, SIQA, HellaSwag, WinoGrande, ARC-e, ARC-c, OBQA) following Hu et al. (2023); Liu et al. (2024). Secondary: WikiText-2 perplexity of the quantized base (Table 14). **Models.** phi-1.5 (1.4B), Llama-3.2-1B, Llama-3.2-3B, Phi-3-mini-4k-instruct (3.8B), Qwen2.5-7B, and Qwen2.5-14B: four architecture families across a 10× size range, chosen so the memory regime shifts from "quantization optional" to "quantization required" within the study.

The two mid-scale models are there for architectural variation, not for another point on the size axis. The rule operates on the list of linear layers, and that list is not the same shape in every family (Table 8, Appendix B). Grouped-query attention shrinks `k_proj` and `v_proj` well below `q_proj`, and Phi-3 fuses its projections into `qkv_proj` and `gate_up_proj`, leaving four kinds of linear layer where Llama has seven. A budget of ρ=.20 therefore buys different things: 4.3% of the weights on Llama-3.2-1B against 16.3% on Phi-3-mini, a factor of 3.8 for the same nominal setting, and a factor of fourteen in the memory it costs. We report this because it is a defect in how we specified the budget. A fraction of *layers* does not transfer across architectures; a fraction of *parameters* would. Comparisons at equal ρ across families are therefore comparisons at unequal memory, and we give Phi-3 a memory-matched budget alongside the nominal one for that reason.

Baselines. Quantized: QLoRA, rsLoRA, QDoRA, LoftQ; fp16 references: LoRA and DoRA. LoftQ runs on three of the six models and fails on Llama-3.2-3B and Qwen2.5-7B for an environment reason, not a property of the method; we do not read its absence in our favour (Appendix I). PiSSA cannot initialize on a quantized base by design (Meng et al., 2024), and LoRA+ and AdaLoRA vary the learning-rate ratio or rank budget, a setting independent of precision, so we leave all three out of this study. Search-based joint allocation (Zhou et al., 2025a; 2026) has no public code; we compare allocation cost analytically, and compare against Unsloth's released selection directly in §7.2. **Hyperparameters.** r=8, LoRA scaling factor 16 (that is, $2r$), lr 3×10^{-4} cosine, 1 epoch, batch 6×8 accumulation, 600 tokens, fp16. **Fidelity.** Each baseline's settings were checked against its official repository; deviations and per-method notes are Appendix I. **Hardware.** One V100-SXM2-32GB per run, a pre-Hopper datacenter GPU. All results in this paper are on this card. We target smaller commodity and edge GPUs, but do not run on them here; §E gives a feasibility analysis with bandwidth-scaled projections, not measurements. **Timing protocol.** Accuracy and timing are measured separately. Accuracy does not depend on load, so accuracy runs share the node with other users; their wall-clock is not a timing measurement, and no table reports it. For timing, each model's methods run in one session on one GPU, *interleaved* so that any drift in machine state affects every method equally, with five or seven repeats per method and 200 measured steps after warmup (60 in the two earliest kept sessions). In twelve of the thirteen sessions, each run starts only once the timing GPU is empty; the earliest phi-1.5 session instead held the node's other GPUs idle for its duration. The node is shared and not reserved. NVML

Table 1: Commonsense-170K 8-task accuracy (%) on phi-1.5, no rsLoRA scaling.

Method	BoolQ	PIQA	SIQA	HellaS	WinoG	ARC-e	ARC-c	OBQA	Avg	Peak GiB	Params
LoRA fp16 (ref)	**65.0**	76.5	76.5	83.2	78.0	80.5	67.1	71.6	74.8	5.18	7.1M
DoRA fp16 (ref)	64.6	76.5	**76.9**	83.1	**78.4**	81.2	**67.6**	71.8	75.0	5.19	7.5M
QLoRA†	64.5	76.6	76.5	83.2	76.9	81.3	66.6	71.9	74.7	**3.83**	7.1M
rsLoRA†	64.7	**77.3**	76.3	**83.5**	77.4	**81.5**	66.5	**74.0**	**75.2**	**3.83**	7.1M
LoftQ‡§	64.6	76.3	76.3	83.4	77.6	**81.5**	67.0	72.1	74.8	**3.83**	7.1M
QDoRA	**65.0**	76.2	76.0	82.9	76.8	81.4	67.2	70.6	74.5	3.84	7.5M
AQLoRA-q†	64.7	76.7	76.2	83.2	77.1	81.2	66.7	72.5	74.8	4.05	7.1M
AQLoRA-s	64.7	74.4	75.4	82.0	75.4	79.0	65.3	69.6	73.2	4.03	6.0M

Avg = mean of the eight task columns; bold = best per column on unrounded values. †Mean of six seeds, §of three (Table 12); unmarked rows are single runs. ‡PEFT one-shot NF4 initialization. Speed is in Tables 3 and 4.

confirms that the timing GPU held its maximum clock (1530 MHz) with no thermal or power throttling, including while other users' jobs ran on other GPUs of the same chassis. Interference can only slow a run, so the samples have a ceiling at the uncontended speed and a slow tail. We therefore estimate speed as the mean of the fastest half of the repeats: three or four runs rather than one lucky maximum, without the interference-slowed tail.

Noise floor and exclusion. Each session measures its own noise floor. Nine sessions include a second, identically configured QLoRA arm, and the floor is the gap between the two QLoRA arms. The four earliest kept sessions predate that arm and use a proxy instead: rsLoRA and LoftQ differ from QLoRA only in a scaling constant and in initialization, so their gap to QLoRA is also a floor. A session is kept when its floor is under 3%; at or above that, the floor is the size of the effect being measured. Five sessions fail, with floors of 3.5 to 8.7%, and are excluded; all five are in the released artifacts. Every model is timed in two or three kept sessions, thirteen in all. Appendix O repeats the analysis under three other estimators, with the excluded sessions included, and with the duplicate-arm sessions alone. The lever ladder (Appendix K) is a single sequential pass and does not follow this protocol. **Allocator.** All runs set PyTorch's `expandable_segments` allocator, applied identically to every method, so it lowers absolute peak memory without affecting any comparison between methods. **Efficiency metrics.** Steps/s and real (non-pad) tokens/s (the timing protocol), plus padding waste, peak and base-model GiB, trainable parameters, and profiling seconds. We treat efficiency as a primary result, next to accuracy. **Artifacts.** Every run writes a `report.json` with its per-task scores, peak memory, wall-clock, and the exact plan used. All of them are released with the code, so each table can be recomputed from the artifacts alone.

5 Main results

On phi-1.5 every method except AQLoRA-s sits in a tight band (74.5 to 75.2 average). Over repeated runs AQLoRA-q (74.78), LoftQ (74.84) and QLoRA (74.69) sit within 0.15 of one another and of the fp16 LoRA reference (74.80); only rsLoRA separates, at 75.16 (Table 12). The gaps among the first four are smaller than the seed spread of any one of them, so the result is parity, not a win. The central rows are repeated because a single seed can overstate a margin, as the next case shows. Llama-3.2-1B makes the point concretely (Table 9). At one seed AQLoRA-q led QLoRA by 1.1 points; over three seeds the margin is 0.3 (72.2 against 71.9), because QLoRA's first seed was its worst. AQLoRA-q is level with LoftQ (72.3, one run) there. Across both models AQLoRA-q matches fp16 LoRA within seed noise and is at or above QLoRA, at a fraction of fp16's memory (base 1.18 against 2.64 GiB on phi-1.5) and higher throughput than QLoRA (Table 4). AQLoRA-s (depth α=.85) trades accuracy for speed. Against QLoRA's multi-seed means it is 1.5 points lower on phi-1.5 (73.2), 1.0 on Llama-3.2-3B (82.25 against 83.25), 0.4 on Phi-3-mini, 0.3 on Llama-3.2-1B (71.6 against 71.9), and 0.1 on Qwen2.5-7B. At seed 42 alone it leads QLoRA on Llama-3.2-1B (71.6 against 71.1), but that is QLoRA's worst seed. The speed setting costs 0.1 to 1.5 points. QDoRA reaches QLoRA accuracy under these settings; under its own recipe it gains 0.89 points and passes AQLoRA-q by

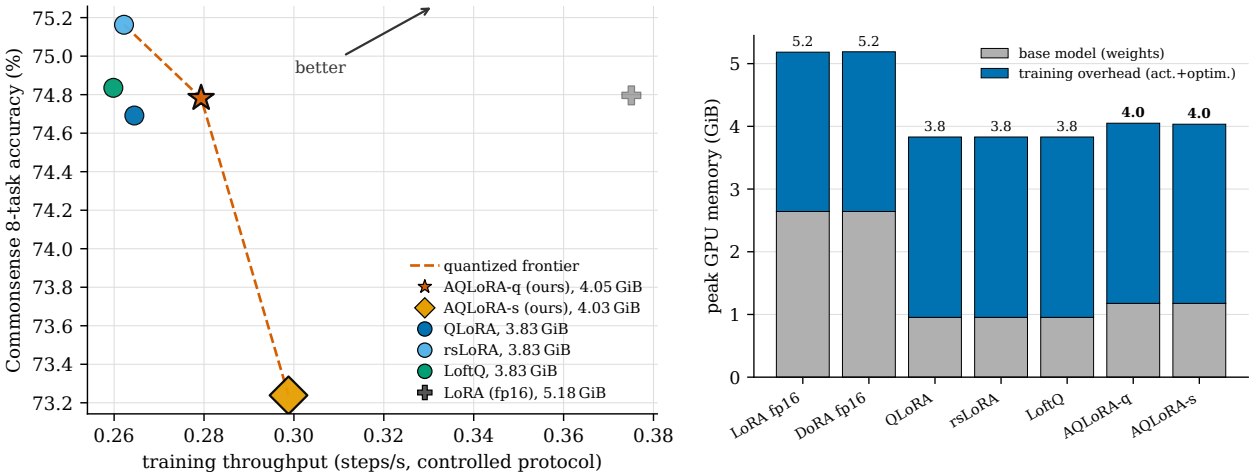


Figure 2: Left: accuracy against training throughput on phi-1.5; accuracy as in Table 1, throughput from Table 3. Right: peak memory.

0.26 (Appendix A). The closest competitor is LoftQ. The two are tied on accuracy: 74.84 for LoftQ over three seeds against 74.78 over six, inside the spread of either. LoftQ uses less memory, 3.83 against 4.05 GiB of peak, since the fp16-protected layers cost 0.22 GiB. AQLoRA-q is faster, by $4.8 \pm 2.0\%$ over QLoRA-class methods across sessions (Table 4). So AQLoRA-q is behind or level on two of the three axes a practitioner weighs, and its case rests on step time alone. A reader who does not care about training time should prefer LoftQ here. Both sit well under fp16 LoRA's 5.18 GiB (Figure 2, right). The accuracy is not the selling point; the speed at parity is (§6).

Table 6 (Appendix B) gives Llama-3.2-3B and Phi-3-mini on the 8-task benchmark, three seeds per cell. Neither changes the accuracy picture. On Llama-3.2-3B the quantized methods sit between 81.7 and 83.5 against an fp16 reference of 84.0, and AQLoRA-q's +0.26 over QLoRA carries p=.32. On Phi-3-mini the three main 4-bit methods span 0.07 points, QLoRA at 87.0, LoftQ at 87.0 and AQLoRA-q at 87.0, with p=.73 against QLoRA. We read the second as the useful one. On a current architecture at 3.8B, the choice among quantized LoRA methods does not move this benchmark. Differences of this size are smaller than the spread between repeated runs of one configuration (§7.4).

The memory column is where the two models separate, and it is the reason we added them. AQLoRA-q costs +0.27 GiB over QLoRA on Llama-3.2-3B and +0.82 on Phi-3-mini for the same ρ, which is the fused-projection effect of Table 8 showing up as measured peak rather than as an estimate. Peak memory is not comparable *across* the two: Llama-3.2 carries a 128K-token vocabulary against Phi-3's 32K, which by itself accounts for most of the 3.8 GiB difference in non-weight memory, and has nothing to do with precision allocation. Only the premium within a model is a property of the method.

5.1 Scale: where quantization stops being a choice

At 1.4B, fp16 LoRA fits everywhere and there is no reason to quantize; the methods above compete in a band best read as parity. Scale changes that. Table 2 runs the core set on Qwen2.5-7B and reports memory at 14B, same protocol throughout. At 7B, fp16 LoRA peaks at 19.1 GiB, which already crowds a 24 GB consumer card, while AQLoRA-q matches its accuracy (89.0 for both) at 10.6 GiB. At 14B, fp16 LoRA is out of reach on commodity hardware: its weights alone are 27.5 GiB, past a 24 GB card, while the quantized methods load the same base in 9 to 11 GiB. In this regime quantization is not an optimization but a requirement, and the question becomes which quantized method to use. AQLoRA does not undercut QLoRA on memory; the fp16-protected layers add 0.5 GiB at 7B and 1.9 GiB at 14B over uniform 4-bit. At 7B that premium buys a measured speedup (Table 3). At 14B, AQLoRA-s reaches slightly higher accuracy than QLoRA (91.3 against 91.1); 14B was not timed under the controlled protocol, so we make no speed claim there. At 7B,

Table 2: Scale study: Qwen2.5-7B accuracy and the 14B memory picture.

Qwen2.5-7B	Avg	Peak GiB
LoRA fp16 (ref)	**89.0**	19.06
AQLoRA-q	89.0	10.55
AQLoRA-s	88.7	10.51
QLoRA	88.8	**10.04**
rsLoRA	87.1	**10.04**

Rows are single runs at seed 42; §7.4 gives three seeds for AQLoRA-q (89.08) and QLoRA (88.80). LoftQ fails on Qwen2.5-7B for an environment reason (Appendix I); DoRA and QDoRA were run on the two small models only. *Qwen2.5-14B memory (32 GB V100).* fp16 LoRA needs 27.5 GiB for the weights alone, past a 24 GB card and near a 32 GB V100's limit; that run was stopped early (43-h ETA) and has no accuracy entry. QLoRA reaches 91.1 at 12.4 GiB, AQLoRA-s 91.3 at 14.2 GiB, and AQLoRA-q 91.1 at 14.3 GiB. The two AQLoRA rows compute the protected layers in fp32, with storage unchanged (§8).

AQLoRA-q is above QLoRA on all three seeds (§7.4). rsLoRA underperforms at the shared learning rate (87.1, §D).

A second task family. We also check that the parity result is not specific to commonsense data. We fine-tune Qwen2.5-7B on a 20K subset of MetaMathQA and score exact-match accuracy on the GSM8K test set (1,319 problems), the MetaMathQA-to-GSM8K protocol of PiSSA (Meng et al., 2024). The same picture holds (Table 7, Appendix B). The three methods land within 0.5 points of one another, inside the ±1-point sampling noise of this test set. The quantized runs peak at 10.05 and 10.55 GiB against fp16's 19.06. The parity claim carries across task families; here too the difference between methods is cost, not accuracy.

6 Where the speed comes from

Table 4 reports speed over the thirteen kept sessions. AQLoRA-s averages 11.3% faster than QLoRA with a between-session standard deviation of 2.3, and it is faster in every session, by 7.0 to 15.2%; the smallest margin exceeds every floor. AQLoRA-q averages 4.8 ± 2.0% and is also faster in every session, by 2.1 to 8.1%. Neither result depends on the estimator or the exclusions (Appendix O). rsLoRA and LoftQ share QLoRA's step cost to within the floor (Table 3), so the same margins hold against them. fp16 LoRA is faster still wherever it fits, by 16.6 to 41.8% over QLoRA in Table 3.

A floor measured inside one session does not predict the spread between sessions. On phi-1.5, Llama-3.2-1B and Qwen2.5-7B, repeated sessions of AQLoRA-s land within 0.1, 0.7 and 0.2 points of one another. On Llama-3.2-3B the two later sessions sit 5.5 and 6.4 points above the first. On Phi-3-mini three kept sessions span 8.2 points, from 7.0 to 15.2, while one of them reports a within-session floor of 1.3%. Excluded sessions differ more: on phi-1.5 they put AQLoRA-s at 16.7 and 6.2%, against 13.0 in a kept session. A speed number from a single session, here or elsewhere, carries an uncertainty that its own error bar does not show. Table 3 is kept beside Table 4 to show the difference.

Per-model means of AQLoRA-s range from 10.5% (Qwen2.5-7B) to 12.9% (phi-1.5). Those two models repeat tightly and differ from each other (Welch $p<.001$). The other three spread by up to 8 points across sessions, and a one-way test across all five models does not separate them ($p=.89$). The gain may depend on the model; with two or three sessions each, these data cannot say how.

For timing on shared hardware this reduces to three rules, each tied to a failure documented in Appendix F. Fix the measurement duration, not the step count, or faster models get noisier numbers. Measure the floor from two runs of the same configuration; a nearly identical method reads too narrow. Repeat the whole session before quoting a margin; a clean floor is necessary, not sufficient. The seed does not control this variance (§7.4).

The lever ladder (Table 16) adds one change at a time on phi-1.5 in a single sequential pass, so its margins are single-session numbers. Length-grouped batching is the largest shared lever, cutting padding waste from

Table 3: Steps/s from one session per model under the controlled protocol.

	QLoRA	rsLoRA	LoftQ	AQLoRA-q	AQLoRA-s	LoRA fp16	floor
phi-1.5[a]	0.265	0.262	0.260	0.279	**0.299**	0.375	1.77%
vs. QLoRA	ref	−0.9%	−1.8%	+5.6%	+13.0%	+41.8%	
Llama-3.2-1B[c]	0.372	0.366	0.371	0.381	**0.414**	0.492	1.78%
vs. QLoRA	ref	−1.8%	−0.5%	+2.4%	+11.1%	+32.1%	
Llama-3.2-3B[c]	0.213	0.210	n/a[b]	0.219	**0.229**	0.265	**2.24%**[d]
vs. QLoRA	ref	−1.8%	n/a[b]	+2.5%	+7.2%	+24.3%	
Phi-3-mini[c]	0.254	0.257	0.257	0.259	**0.271**	0.296	1.33%
vs. QLoRA	ref	+1.3%	+1.2%	+2.1%	+7.0%	+16.6%	
Qwen2.5-7B[a]	0.194	0.194	n/a[b]	0.201	**0.215**	0.237	0.10%
vs. QLoRA	ref	−0.1%	n/a[b]	+3.6%	+10.5%	+21.9%	

Seven interleaved repeats per method; estimate = mean of the fastest half. [a]60 measured steps; [c]200 (§4). [b]LoftQ unavailable (Appendix I); floor rests on rsLoRA alone. [d]Duplicate-arm floor; the other four are proxy floors and read narrow (§4). Percentages are comparable within a row only. The phi-1.5 AQLoRA arms carry rsLoRA scaling, a constant factor with no compute cost; rsLoRA itself, which differs from QLoRA only by that factor, lands 0.9% away.

Table 4: Speed over independent sessions, percent faster than QLoRA.

Model	clean sessions	AQLoRA-q	AQLoRA-s
phi-1.5	2	$+6.3 \pm 0.9$	$+12.9 \pm 0.0$
Llama-3.2-1B	2	$+5.1 \pm 3.9$	$+11.4 \pm 0.5$
Llama-3.2-3B	3	$+4.4 \pm 1.7$	$+11.1 \pm 3.5$
Phi-3-mini	3	$+4.8 \pm 3.0$	$+11.1 \pm 4.1$
Qwen2.5-7B	3	$+4.0 \pm 0.4$	$+10.5 \pm 0.1$
pooled	13	$+4.8 \pm 2.0$	$+11.3 \pm 2.3$

Sessions with a floor of 3% or more are excluded (five; §4).

39.8% to 3.7%. Truncating the backward pass at the lowest adapted block adds 5.6% at α=.7 (0.287 to 0.303 steps/s). With gradient checkpointing off, depth placement at α=.7 runs 42% faster than QLoRA at 23% less memory, because stopping the backward pass early also drops the lower-block activations.

On accuracy against throughput alone, fp16 LoRA and rsLoRA form the frontier (Figure 2). Within the 4-bit memory class, both settings of ours are on the frontier and QLoRA and LoftQ are not, though the accuracy gaps are inside seed noise. The figure labels each point with its peak memory, since the frontier is only meaningful within a class.

7 Ablations

7.1 Which layers to adapt

Picking the layers to adapt by weight density does not help. We test it at two budgets ($\alpha \in \{0.5, 0.7\}$) against a weight-norm ranking (the Act-LoRA baseline) and a random subset (Table 15). The table has two columns that differ in seed and in rsLoRA scaling, which the seed-42 runs used; comparisons hold within a column. Random is ahead of density in all four cells, by 1.0 and 1.1 points at α=.5 and by 0.2 and 0.5 at α=.7, and density is last of the three in every cell. The density score gives no benefit and we do not use it: the quality setting adapts every layer.

What matters more is the budget. In the seed-42 column, dropping from every layer to α=.7 costs 0.7 points (75.56 to 74.84). In the lever ladder, cutting to α=.7 with depth placement but no backward truncation raises throughput from 0.260 to 0.287 steps/s (Table 16), so most of that gain comes from the adapter count. Restricting the adapters to the top blocks (depth) costs more accuracy at the same budget (71.2 at α=.7);

in exchange, truncating the backward pass adds 5.6% throughput. At α=.85 the depth rule scores 74.6 in that column. Without rsLoRA scaling, which is the shipped speed setting, it scores 73.2 (Table 1).

7.2 Precision budget

Base perplexity improves with the protection budget under both NF4 and the harsher fp4 (Appendix G prices this). That perplexity gain does not carry to the downstream task. The budget ρ=.20 was chosen on this evidence. Table 18 then sweeps it end to end and reports 8-task accuracy, memory, and step time. Accuracy shows no trend: the single-seed budgets span 0.6 points, less than the 0.76 that seeds 42 to 44 span at ρ=.20 alone (74.35 to 75.11). Base perplexity keeps falling where accuracy does not: the residual tier of Appendix H at ρ=.40 reaches 18.83, against 18.94 for fp16 at ρ=.20, and scores 74.96. Base perplexity, which we used to choose the budget, is a poor guide to it.

Since accuracy is flat here, the budget should be set by cost. At ρ=.05 the premium over QLoRA falls to 0.06 GiB and in one session the step-time advantage is 4.1% against a 2.8% noise floor. The flat result does not hold on Phi-3-mini, where ρ=.05 scores lower, so it is not a default (Appendix M).

Does the ranking select better than chance? On phi-1.5, accuracy does not respond to how many layers are protected. The next question is whether it responds to *which* ones. The speed side cannot answer it, because any K fp16 layers skip the same dequantization. We hold the protected memory fixed and vary only the selection: top-K by NF4 error, the K least damaged layers, and ten random draws matched to top-K. Base perplexity needs no training and is deterministic given the layer set, so these numbers are exact.

Matching the draws by layer count, as our first control did, is not enough. Where layer sizes vary, the most damaged layers are smaller than average (5.5M against 8.4M parameters on phi-1.5), so a random draw of the same count carries more parameters. On Llama-3.2-1B the draws held up to 11% more memory than top-K and won on perplexity for that reason alone. We therefore match on parameter total and let the count float. That holds the protected memory equal to within 0.1% in every setting reported here, and it is the control we recommend for any budgeted allocation rule.

Under that control the ranking has no measurable advantage (Appendix N lists every setting). Across five models and three budgets, top-K beats the median matched draw in 11 of 15 settings, and 50 of the 150 draws beat top-K where chance predicts 75. A sign test does not separate that from chance (p=.12), and the median gap is 0.01 perplexity. The direction is mildly favourable and the size is negligible.

Downstream the selections are indistinguishable on both models we tested. On phi-1.5 at seed 42, top-K and matched random score 74.88 and 74.63, against a seed spread near 0.25 (Table 12); bottom-K, which protects 79% more parameters, scores 74.89. On Phi-3-mini, where the perplexity gap is widest, top-K and matched random reach 87.04 and 87.03 over three seeds. Phi-3-mini is also the model on which no quantized method separates (§5), so phi-1.5 is the more discriminating test, at one seed.

What protection buys, speed and lower base perplexity, therefore comes from protecting some layers, not from protecting the right ones. NF4 reconstruction error is not doing the work we attributed to it. This does not touch the speed result, which depends only on how many layers avoid dequantization, nor the accuracy parity. It does remove the mechanism we offered: that protecting the most damaged layers is what preserves quality. A reader should treat the rule as a cheap and reproducible way to fix a layer count, not as a sensitivity criterion shown to work. A parameter-matched random selection does as well; the ranking costs 5 to 37 s of CPU time and gives a fixed, documented choice, which is the only reason it remains the default. The test is at whole-layer granularity, so it says nothing about the column-granularity selection of QEFT and OWQ.

A cheaper protection tier, a second NF4 stage over each protected layer's own residual, also works and is also not worth using; Appendix H reports it as a negative result.

We also compare our rule against the only deployed alternative we know of: Unsloth's released dynamic-4bit selection for Qwen2.5-7B, a hand-curated list covering 34 of the 196 linear layers. Substituting their list for ours in an otherwise identical run gives 89.0 on the 8-task average, matching our 89.0. The two selections are not the same layers: their list names 10 modules, ours 30, and only 6 are shared. Two different selections at

different budgets landing on the same accuracy is the ranking control (Appendix N) seen from the other side. The practical point survives it: the budgeted rule needs no per-model curation to match the hand-tuned selection.

The rank sweep, the rsLoRA composition, and the edge estimate are in Appendices C, D, and E; the method ordering holds at every rank, and the composition does not reach significance.

7.3 Why the two rules fail

The density and weight-norm rules pick almost the same layers: 68 of 72 at α=.5 and 99 of 100 at α=.7. The density score is therefore close to a weight-magnitude ranking, and where the two sets differ by one layer they score within 0.1 to 0.3 points of each other. The sets differ from a random draw in two ways. They sit higher in the network: their mean adapted block is 12.5 to 13.4 of 0 to 23, against 10.8 to 11.7 for random draws (Table 15). They also skew by module type: at α=.5 the density rule adapts all 24 query projections and one of 24 value projections, where a random draw takes 9 to 12 of each.

To separate the two, we drew random layers with the density rule's count in each third of the network. The depth profile then matches, while the layers do not (36 to 40 of 72 shared). In both columns of Table 15 this depth-matched control lands 0.25 points below unconstrained random and 0.8 to 0.9 above density: 74.3 against 74.6 and 73.6 with rsLoRA scaling, 73.3 against 73.5 and 72.4 without. Depth accounts for 23 to 25% of the density rule's loss. The rest comes from which layers it picks; module type is the likely cause, since the depth-matched sets spread across types as random does, but we did not test it separately.

The depth rule is the extreme case of moving adapters up. At α=.7 its mean block is 15.2 and it scores 3.7 points below density at the same budget. In exchange, no layer below its lowest adapted block needs a gradient, so the backward pass stops there. The density rule moves adapters up without emptying the lowest blocks, so it pays an accuracy cost and gets no backward-pass cut.

The error ranking fails for two reasons that the corrected control separates. First, NF4 error falls with layer size ($r = -0.19$): the 29 most damaged layers on phi-1.5 average 5.5M parameters against 8.4M overall. A control matched on count therefore hands the random draws more memory, which is the artifact of §7.2. Second, the base-perplexity advantage that survives a fair control is at most 0.05 against the median matched draw, and it does not reach accuracy. That is consistent with the premise of LoftQ (Li et al., 2024), that low-rank adapters compensate quantization error: a difference this small between two protected sets is within what the adapter absorbs. Protection pays in speed, by removing dequantization work, and lowers base perplexity as the budget grows (Appendix G). Which layers carry the fp16 does not measurably matter once the model is fine-tuned.

7.4 Seeds

On Qwen2.5-7B we reran the central pair on three seeds: AQLoRA-q reaches 89.08 ± 0.19 against QLoRA's 88.80 ± 0.11, above it on every seed with no overlap (worst AQLoRA-q seed 88.97, best QLoRA seed 88.91). A gap of 0.28 points is why we call the accuracy parity; its sign is the same on every seed.

On phi-1.5 the same pair ties (Table 12): AQLoRA-q averages 74.78 ± 0.25 against QLoRA's 74.69 ± 0.22 over the six runs of Table 12, so the seed-42 gap (74.88 against 74.74) does not survive repetition. At three seeds AQLoRA-q also looked far less reproducible (F=71, two-sided p=.03); a fourth seed per arm brought that to F=1.9 (p=.60), and six seeds to F=1.2. A variance ratio from three runs per arm can flip on one draw.

These error bars do not measure seed variation. Three runs of AQLoRA-q at the *same* seed spread as widely as six runs at different seeds (sd 0.24 against 0.25). Our training does not enable deterministic kernels, so the error bars measure run-to-run nondeterminism, and fixing the seed does not make results reproducible.

8 Limitations

The largest limitation is the selection rule. It ranks layers by NF4 reconstruction error. Tested against parameter-matched random selections on five models at three budgets, it is not measurably ahead, and the advantage it does show on base perplexity does not reach downstream accuracy (§7.2). Protecting layers buys speed and lowers base perplexity; protecting the ones our criterion picks is not measurably better than protecting an arbitrary set of the same size. The speed result is unaffected, since it depends on how many layers avoid dequantization rather than which. The mechanism offered for the recipe is not supported by the control. A reader should take AQLoRA as a reproducible way, costing 5 to 37 s, to fix how many layers are protected, not as a demonstration that sensitivity ranking works at whole-layer granularity. The test covers five models at three budgets each, ten draws per setting. That is enough to say the criterion earns nothing measurable at whole-layer granularity. It does not rule out a gain at finer granularity, which we did not test.

At 14B some of the layers our rule protects overflow in fp16 backward on pre-bf16 GPUs. This silently collapses the AMP loss scaler while the training loss still looks plausible. Uniform QLoRA is immune, and our mitigation computes the protected layers in fp32 at unchanged storage. On bf16 hardware the issue does not arise. We do not expect AQLoRA to win at matched memory. Paying for the fp16-protected layers by shrinking the batch or sequence trades the memory saving back, so the gain is speed, not accuracy per byte. The speed setting gives up 0.1 to 1.5 points of accuracy against QLoRA (§5), so its speed is not free. Our speedups are on the stock stack and pre-Hopper or edge GPUs; kernel-level and FP8 methods are faster where their hardware applies, and while AQLoRA's levers should compose with them, we do not measure that here. The phone numbers are projections from prior measured data, not on-device runs. Our headline tables fix one epoch, and the rank sweep (§C) covers $r \in \{4, 8, 16, 32\}$ on phi-1.5 and $r \in \{8, 16, 32\}$ on Llama-3.2-1B, not the larger models.

9 Conclusion

AQLoRA sets per-layer precision and adapter placement for quantized LoRA from one CPU pass over the weights, with no search. Across six models and four architecture families its accuracy matches QLoRA's and trails fp16 LoRA by less than a point, and across thirteen timing sessions its speed setting trains $11.3 \pm 2.3\%$ faster than well-tuned QLoRA, for 0.04 to 0.8 GiB more peak memory on the timed models. The result is speed at parity, not higher accuracy. No single part is new: protection by sensitivity, selective placement, and early backward stopping all have prior work, which we cite. What we add is a recipe that composes them from one data-free pass, and the controls that show which part of it does the work. It matches the accuracy of a hand-curated industrial selection where search-based allocators need repeated calibration passes. It also yields three measurement findings. How many layers are protected sets the speed; which ones does not measurably change accuracy. A noise floor measured inside one session understates the spread between sessions: on Phi-3-mini, a 1.3% floor against an 8.2-point span. The seed controls almost none of the run-to-run variance. The speed result holds on the stock stack, on hardware where kernel- and FP8-based methods do not run. A measured phone-GPU system remains future work.

Broader Impact Statement

AQLoRA lowers the time and hardware needed to fine-tune a language model on older and smaller GPUs; the main benefit is access for groups without current datacenter hardware, and faster training means less energy per run. The same access cuts both ways, since cheaper fine-tuning is cheaper for harmful adaptation too. We see no risk specific to this work beyond those that apply to efficient fine-tuning in general.

References

Kanghyun Choi, Hyeyoon Lee, SunJong Park, Dain Kwon, and Jinho Lee. FALQON: Accelerating LoRA fine-tuning with low-bit floating-point arithmetic. In *NeurIPS*, 2025. arXiv:2510.24061.

Yanxia Deng, Aozhong Zhang, Selcuk Gurses, Naigang Wang, Zi Yang, and Penghang Yin. CLoQ: Enhancing fine-tuning of quantized LLMs via calibrated LoRA initialization. *arXiv preprint arXiv:2501.18475*, 2025.

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. In *NeurIPS*, 2023. arXiv:2305.14314.

Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. HAWQ-V2: Hessian aware trace-weighted quantization of neural networks. In *NeurIPS*, 2020. arXiv:1911.03852.

Razvan-Gabriel Dumitru, Vikas Yadav, Rishabh Maheshwary, Paul-Ioan Clotan, Sathwik Tejaswi Madhusudhan, and Mihai Surdeanu. Layer-wise quantization: A pragmatic and effective method for quantizing LLMs beyond integer bit-levels. *arXiv preprint arXiv:2406.17415*, 2024.

Jiaxiang Geng, Lunyu Zhao, Yiyi Lu, and Bing Luo. MobileFineTuner: A mobile-native framework for on-device LLM fine-tuning in real-world embedded AI applications. *arXiv preprint arXiv:2512.08211*, 2025.

Han Guo, Philip Greengard, Eric P Xing, and Yoon Kim. LQ-LoRA: Low-rank plus quantized matrix decomposition for efficient language model finetuning. In *ICLR*, 2024. arXiv:2311.12023.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *ICLR*, 2022.

Zhiqiang Hu, Lei Wang, Yihuai Lan, Wanyu Xu, Ee-Peng Lim, Lidong Bing, Xing Xu, Soujanya Poria, and Roy Ka-Wei Lee. LLM-Adapters: An adapter family for parameter-efficient fine-tuning of large language models. In *EMNLP*, 2023. Commonsense-170K protocol; arXiv:2304.01933, ACL Anthology 2023.emnlp-main.319.

Kai Huang, Hanyun Yin, Heng Huang, and Wei Gao. Towards green AI in fine-tuning large language models via adaptive backpropagation. In *ICLR*, 2024. GreenTrainer, arXiv:2309.13192.

Wei Huang, Yi Ge, Shuai Yang, Yicheng Xiao, Huizi Mao, Yujun Lin, Hanrong Ye, Sifei Liu, Ka Chun Cheung, Hongxu Yin, Yao Lu, Xiaojuan Qi, Song Han, and Yukang Chen. QeRL: Beyond efficiency – quantization-enhanced reinforcement learning for LLMs. *arXiv preprint arXiv:2510.11696*, 2025.

Damjan Kalajdzievski. A rank stabilization scaling factor for fine-tuning with LoRA. *arXiv preprint arXiv:2312.03732*, 2023.

Samir Khaki, Xiuyu Li, Junxian Guo, Ligeng Zhu, Chenfeng Xu, Konstantinos N Plataniotis, Amir Yazdanbakhsh, Kurt Keutzer, Song Han, and Zhijian Liu. SparseLoRA: Accelerating LLM fine-tuning with contextual sparsity. In *ICML*, 2025. arXiv:2506.16500.

Seokmin Ko. LoRA is slower than you think. *arXiv preprint arXiv:2507.08833*, 2025.

Changhun Lee, Jun-gyu Jin, Younghyun Cho, and Eunhyeok Park. QEFT: Quantization for efficient fine-tuning of LLMs. In *Findings of EMNLP*, 2024a. arXiv:2410.08661.

Changhun Lee, Jungyu Jin, Taesu Kim, Hyungjun Kim, and Eunhyeok Park. OWQ: Outlier-aware weight quantization for efficient fine-tuning and inference of large language models. In *AAAI*, 2024b. arXiv:2306.02272.

Yixiao Li, Yifan Yu, Chen Liang, Pengcheng He, Nikos Karampatziakis, Weizhu Chen, and Tuo Zhao. LoftQ: LoRA-fine-tuning-aware quantization for large language models. In *ICLR*, 2024. arXiv:2310.08659.

Baohao Liao, Yan Meng, and Christof Monz. Parameter-efficient fine-tuning without introducing new latency. In *ACL*, 2023. PaFi, arXiv:2305.16742.

Baohao Liao, Christian Herold, Shahram Khadivi, and Christof Monz. ApiQ: Finetuning of 2-bit quantized large language model. In *EMNLP*, 2024. arXiv:2402.05147.

Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-Ting Cheng, and Min-Hung Chen. DoRA: Weight-decomposed low-rank adaptation. In *ICML*, 2024. arXiv:2402.09353.

Fanxu Meng, Zhaohui Wang, and Muhan Zhang. PiSSA: Principal singular values and singular vectors adaptation of large language models. In *NeurIPS*, 2024. arXiv:2404.02948.

Rui Pan, Xiang Liu, Shizhe Diao, Renjie Pi, Jipeng Zhang, Chi Han, and Tong Zhang. LISA: Layer-wise importance sampling for memory-efficient large language model fine-tuning. In *NeurIPS*, 2024. arXiv:2403.17919.

Juneyoung Park, Eunbeen Yoon, Seongwan Kim, and Jaeho Lee. LCSB: Layer-cyclic selective backpropagation for memory-efficient on-device LLM fine-tuning. *arXiv preprint arXiv:2602.13073*, 2026.

Haotong Qin, Xudong Ma, Xingyu Zheng, Xiaoyang Li, Yang Zhang, Shouda Liu, Jie Luo, Xianglong Liu, and Michele Magno. Accurate LoRA-finetuning quantization of LLMs via information retention. In *ICML*, 2024. IR-QLoRA, arXiv:2402.05445.

Andreas Rücklé, Gregor Geigle, Max Glockner, Tilman Beck, Jonas Pfeiffer, Nils Reimers, and Iryna Gurevych. AdapterDrop: On the efficiency of adapters in transformers. In *EMNLP*, 2021. arXiv:2010.11918.

Abdulmalek Saket. Aletheia: Gradient-guided layer selection for efficient LoRA fine-tuning across architectures. *arXiv preprint arXiv:2604.15351*, 2026.

Congzheng Song and Xinyu Tang. Memory-efficient backpropagation for fine-tuning LLMs on resource-constrained mobile devices. In *EMNLP (Industry Track)*, 2025. MeBP, Apple; arXiv:2510.03425.

Yi-Lin Sung, Jaemin Cho, and Mohit Bansal. LST: Ladder side-tuning for parameter and memory efficient transfer learning. In *NeurIPS*, 2022. arXiv:2206.06522.

Qitao Tan, Xiaoying Song, Jin Lu, Guoming Li, Jun Liu, Lingzi Hong, Caiwen Ding, Jundong Li, Xiaoming Zhai, Shaoyi Huang, Wei Niu, and Geng Yuan. End-to-end on-device quantization-aware training for LLMs at inference cost. *arXiv preprint arXiv:2509.00031*, 2025.

Tether AI. An edge-first generalized LLM LoRA fine-tuning framework for heterogeneous GPUs. https://huggingface.co/blog/qvac/fabric-llm-finetune, 2025.

Unsloth team. Unsloth dynamic 4-bit quantization. https://unsloth.ai/blog/dynamic-4bit, 2024. Blog post and released model collection, December 2024.

Yiding Wang, Decang Sun, Kai Chen, Fan Lai, and Mosharaf Chowdhury. Egeria: Efficient DNN training with knowledge-guided layer freezing. In *EuroSys*, 2023. arXiv:2201.06227.

Sunghyeon Woo, Baeseong Park, Byeongwook Kim, Minjung Jo, Se Jung Kwon, Dongsuk Jeon, and Dongsoo Lee. DropBP: Accelerating fine-tuning of large language models by dropping backward propagation. In *NeurIPS*, 2024. arXiv:2402.17812.

Hao Wu, Patrick Judd, Xiaojie Zhang, Mikhail Isaev, and Paulius Micikevicius. Integer quantization for deep learning inference: Principles and empirical evaluation. *arXiv preprint arXiv:2004.09602*, 2020.

Yuhui Xu, Lingxi Xie, Xiaotao Gu, Xin Chen, Heng Chang, Hengheng Zhang, Zhengsu Chen, Xiaopeng Zhang, and Qi Tian. QA-LoRA: Quantization-aware low-rank adaptation of large language models. In *ICLR*, 2024. arXiv:2309.14717.

Junjie Yin, Jiahao Dong, Yingheng Wang, Christopher De Sa, and Volodymyr Kuleshov. ModuLoRA: Finetuning 2-bit LLMs on consumer GPUs by integrating with modular quantizers. In *TMLR*, 2024. arXiv:2309.16119; TMLR Featured Certification.

Zhongzhi Yu, Zheng Wang, Yuhan Li, Haoran You, Ruijie Gao, Xiaoya Zhou, Sreenidhi Reedy Bommu, Yang Katie Zhao, and Yingyan Celine Lin. EDGE-LLM: Enabling efficient large language model adaptation on edge devices via layerwise unified compression and adaptive layer tuning and voting. In *DAC*, 2024. arXiv:2406.15758.

Cheng Zhang, Jeffrey T H Wong, Can Xiao, George A Constantinides, and Yiren Zhao. QERA: an analytical framework for quantization error reconstruction. In *ICLR*, 2025a. arXiv:2410.06040.

Chuyan Zhang, Kefan Wang, and Yun Gu. Beyond low-rank tuning: Model prior-guided rank allocation for effective transfer in low-data and large-gap regimes. In *ICCV*, 2025b. SR-LoRA; arXiv:2507.00327.

Qingru Zhang, Minshuo Chen, Alexander Bukharin, Pengcheng He, Yu Cheng, Weizhu Chen, and Tuo Zhao. Adaptive budget allocation for parameter-efficient fine-tuning. In *ICLR*, 2023. AdaLoRA, arXiv:2303.10512.

Changhai Zhou, Shiyang Zhang, Yuhua Zhou, Qian Qiao, Jun Gao, Shichao Weng, Weizhong Zhang, and Cheng Jin. Balancing fidelity and plasticity: Aligning mixed-precision fine-tuning with linguistic hierarchies. *arXiv preprint arXiv:2505.03802*, 2025a. QR-Adaptor; earlier versions titled "Efficient Fine-Tuning of Quantized Models via Adaptive Rank and Bitwidth".

Changhai Zhou, Shiyang Zhang, Yuhua Zhou, Qian Qiao, Jun Gao, Cheng Jin, Kaizhou Qin, and Weizhong Zhang. AutoQRA: Joint optimization of mixed-precision quantization and low-rank adapters for efficient LLM fine-tuning. *arXiv preprint arXiv:2602.22268*, 2026.

Hongyun Zhou, Xiangyu Lu, Wang Xu, Conghui Zhu, Tiejun Zhao, and Muyun Yang. LoRA-drop: Efficient LoRA parameter pruning based on output evaluation. In *COLING*, 2025b. arXiv:2402.07721.

Zikai Zhou, Qizheng Zhang, Hermann Kumbong, and Kunle Olukotun. LowRA: Accurate and efficient LoRA fine-tuning of LLMs under 2 bits. In *ICML*, 2025c. arXiv:2502.08141.

Ligeng Zhu, Lanxiang Hu, Ji Lin, Wei-Chen Wang, Wei-Ming Chen, Chuang Gan, and Song Han. PockEngine: Sparse and efficient fine-tuning in a pocket. In *MICRO*, 2023. arXiv:2310.17752.

Zhanda Zhu, Qidong Su, Yaoyao Ding, Kevin Song, Shang Wang, and Gennady Pekhimenko. LoRAFusion: Efficient LoRA fine-tuning for LLMs. In *EuroSys*, 2026. arXiv:2510.00206.

A On the shared levers and the ranking

Table 1 gives every method the same shared settings (fused AdamW, dropout 0, length-grouped batching), which do not match each method's own published recipe. To test whether that choice drives the result, we reran the four contested methods on phi-1.5 under their native settings instead: paged 8-bit optimizer, dropout 0.05, no length grouping, everything else unchanged (Table 5).

Every method scores higher under its native recipe, so the shared levers cost each of them accuracy: 0.27 to 0.89 points. For the three methods that share a compute path (QLoRA, LoftQ, AQLoRA-q) the shift is small and nearly uniform, +0.27 to +0.40 points, and their order is unchanged. QDoRA is the exception: it gains +0.89, more than twice as much, and moves from last under the shared levers to second under its own recipe, passing AQLoRA-q by 0.26 points. So the shared levers are not neutral for every method, and a reader should not read Table 1 as a recipe-independent ranking. AQLoRA-q stays at or above QLoRA in both settings, which is the comparison it rests on. QDoRA was not timed under the controlled protocol, so we make no speed comparison with it.

B Additional accuracy tables and figures

C Rank sweep

The results so far fix r=8. Table 10 repeats the quantized set at $r \in \{4, 8, 16, 32\}$ on phi-1.5, same protocol throughout. Two things hold across the sweep and one baseline breaks.

AQLoRA-q is at or above QLoRA at all four ranks, though never by much: +0.26 at r=16, +0.03 at r=8 over three seeds, and +0.01 at r=32. The parity result is therefore not an artifact of one rank, and neither

Table 5: Native recipe against the shared levers, phi-1.5 (8-task avg, %).

Method	native	shared levers	change
LoftQ	**75.53**	**75.13**	+0.40
QDoRA	75.41	74.52	+0.89
AQLoRA-q	75.15	74.88	+0.27
QLoRA	75.08	74.74	+0.34

Both columns are seed 42, so the pairing holds within each row; the shared-lever column is a single run, not the multi-seed mean of Table 1.

Table 6: Llama-3.2-3B and Phi-3-mini, 8-task accuracy (%), mean of three seeds.

Method	BoolQ	PIQA	SIQA	HellaS	WinoG	ARC-e	ARC-c	OBQA	Avg	Peak GiB	Params
Llama-3.2-3B											
LoRA fp16 (ref)	**71.6**	**86.4**	**80.8**	**95.1**	**86.4**	**89.5**	**77.7**	**84.1**	**84.0**	11.96	12.2M
QLoRA	70.8	85.8	80.5	94.7	84.9	89.2	77.4	82.6	83.3	**8.07**	12.2M
rsLoRA	69.2	83.8	79.8	93.3	83.3	87.0	74.6	82.9	81.7	**8.07**	12.2M
AQLoRA-q	70.6	86.1	80.7	94.8	85.6	89.1	77.4	83.9	83.5	8.33	12.2M
AQLoRA-s	69.1	85.0	80.3	94.4	85.2	87.6	75.9	80.5	82.2	8.30	10.3M
Phi-3-mini											
LoRA fp16 (ref)	71.3	**88.4**	81.8	**94.9**	**86.3**	**96.2**	**88.9**	**90.9**	**87.3**	9.31	12.6M
QLoRA	70.9	88.1	**81.9**	94.6	85.8	95.8	88.5	90.1	87.0	**4.30**	12.6M
LoftQ	70.8	88.3	81.8	94.7	86.0	96.1	88.4	89.5	87.0	4.34	12.6M
rsLoRA	**71.7**	88.1	81.7	94.6	85.9	95.8	87.3	89.8	86.9	**4.30**	12.6M
AQLoRA-q	71.0	88.3	81.5	94.7	85.7	96.1	88.2	90.8	87.0	5.12	12.6M
AQLoRA-s	70.3	87.5	81.1	94.5	85.6	95.8	88.0	89.7	86.6	5.09	10.6M

LoftQ is absent on Llama-3.2-3B for the environment reason given above; DoRA and QDoRA were run on phi-1.5 and Llama-3.2-1B only. Peak memory is comparable within a model, not between the two: the vocabulary sizes differ by 4×.

Table 7: GSM8K accuracy (%) after MetaMathQA-20K fine-tuning, Qwen2.5-7B. Three methods were run on this task, one run each.

Method	GSM8K	Peak GiB
QLoRA	**84.9**	**10.05**
LoRA fp16 (ref)	84.8	19.06
AQLoRA-q	84.4	10.55

is its size. Its memory premium over QLoRA is constant at 0.22 GiB, which is what the rule predicts: the fp16-protected set is fixed, while the adapters grow with rank, so the premium falls from 5.7% of peak at r=8 to 5.3% at r=32. LoftQ matches AQLoRA-q within noise at all four ranks, differing by at most 0.1 points and in both directions; we read that as a tie, not a win for either, and note that LoftQ reaches it at 0.22 GiB less peak memory. Rank itself matters more than the choice among these methods: the working methods gain 1.7 to 1.9 points from r=4 to r=32, a larger effect than any gap between them at a fixed rank.

Table 8: What the allocation rule sees in each family.

Model	linear layers	layer kinds	size CV	% params at ρ=.20	premium GiB
phi-1.5	144	6	0.71	13.2%	0.220
Llama-3.2-1B	112	7	0.82	**4.3%**	**0.058**
Llama-3.2-3B	196	7	0.67	6.8%	0.265
Phi-3-mini (3.8B)	128	4	0.52	16.3%	0.817

Size CV is the coefficient of variation of layer parameter counts; premium is fp16 storage for the top-K set over all-NF4.

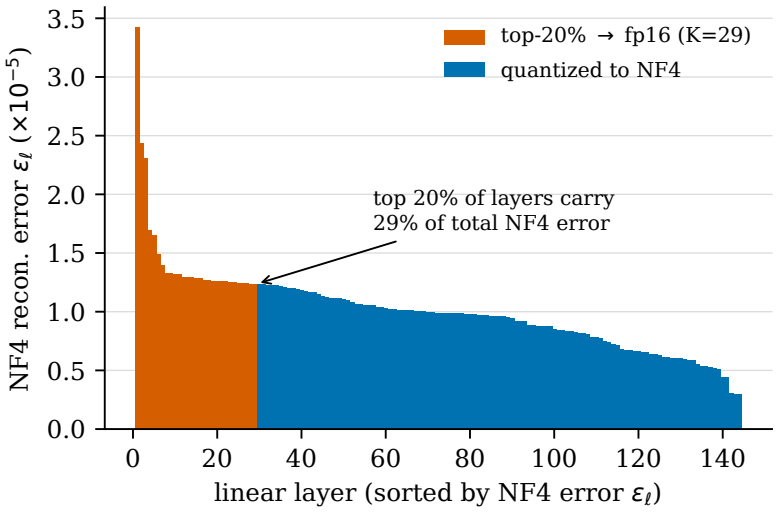


Figure 3: Sorted per-layer NF4 error on phi-1.5. Orange: layers kept in fp16.

Table 9: Commonsense-170K 8-task accuracy (%) on Llama-3.2-1B, no rsLoRA scaling.

Method	BoolQ	PIQA	SIQA	HellaS	WinoG	ARC-e	ARC-c	OBQA	Avg	Peak GiB	Params
LoRA fp16 (ref)	63.8	78.2	74.6	**87.6**	**71.8**	**77.3**	59.4	**72.4**	**73.1**	7.81	5.6M
DoRA fp16 (ref)	**64.9**	**78.5**	73.9	86.8	71.0	76.8	59.6	70.2	72.7	7.82	6.0M
QLoRA†	63.2	77.9	74.2	85.9	70.8	75.6	58.9	68.7	71.9	**6.47**	5.6M
LoftQ	63.5	77.6	74.6	86.2	71.6	75.6	59.1	69.8	72.3	6.58	5.6M
QDoRA	63.2	78.4	74.1	85.9	70.8	75.3	**59.9**	69.8	72.2	6.68	6.0M
AQLoRA-q†	63.8	76.9	**74.7**	86.4	70.8	75.8	59.0	70.1	72.2	6.53	5.6M
AQLoRA-s	63.7	77.0	73.7	85.8	70.2	75.1	59.3	68.4	71.6	6.51	4.8M

Conventions follow Table 1. †Mean of three seeds; unmarked rows are single runs. rsLoRA is omitted: on this model every rsLoRA-scaled run fell to 65 to 67 at the shared learning rate (rsLoRA itself 66.6; Appendix D).

That sweep is phi-1.5 only, so we repeated $r \in \{16, 32\}$ on Llama-3.2-1B (Table 11, one seed per cell). Parity reproduces at r=16: AQLoRA-q, QLoRA and LoftQ land within 0.5 points, with AQLoRA-q 0.33 below QLoRA, inside QLoRA's seed spread on this model (0.71, Table 9). Two things differ from phi-1.5. Rank does not help here, since every method is lower at r=16 than at r=8, the opposite of the phi-1.5 trend. And at r=32 QLoRA and AQLoRA-q both fall to chance while LoftQ degrades but survives at 63.2. The collapse is inherited from the baseline rather than caused by the protection scheme, and it has the same origin as the rsLoRA case: the runs finish with normal eval loss at a rate tuned for r=8. We therefore scope the rank conclusions to phi-1.5 and treat r=32 on Llama-3.2-1B as unusable at the shared learning rate. LoftQ's robustness there is real and we do not explain it.

The exception is rsLoRA, which is the strongest method at r=4 (74.78) and r=8 (75.16 over six seeds), then degrades at r=16 (73.86) and collapses to chance at r=32 (37.47; HellaSwag 25.7, ARC-c 23.8 against a 25 floor). The cause is the scaling, not the method. rsLoRA divides the LoRA scaling factor by $\sqrt{r}$ instead of by r, so with our factor of $2r$ the update is 2.8× larger at r=8, 4× at r=16, and 5.7× at r=32. At our shared learning rate the largest of these overshoots. This is the same failure we report on Llama-3.2-1B (§D), here reached by raising the rank instead of changing the model. Two practical notes follow. A method that wins at one rank under a shared learning rate can lose badly at another, so single-rank comparisons are fragile. And the collapse is invisible in the training signal: the r=32 run ends at eval loss 0.108, close to its r=16 value of 0.103, while its task accuracy is at chance. Held-out loss is not a safe proxy for downstream accuracy here.

Table 10: Rank sweep on phi-1.5 (8-task avg, %).

Method	r=4	r=8	r=16	r=32
AQLoRA-q	74.2	74.8	**75.5**	75.9
QLoRA	74.1	74.8	75.2	75.9
rsLoRA	**74.8**	**75.3**	73.9	37.5*
LoftQ	74.1	74.8	75.5	**76.0**

*Chance-level collapse from rsLoRA scaling at the shared learning rate, not a crash: the run finished with a normal eval loss (0.108). The r=8 column is the mean of seeds 42 to 44 (Table 12; Table 1 extends three of these rows to six seeds); the other ranks are seed 42.

Table 11: Rank on Llama-3.2-1B, 8-task average (%), one seed except r=8, which is the multi-seed mean of Table 9 where available.

Method	r=8	r=16	r=32
LoftQ	**72.27**	**70.06**	**63.20**
QLoRA	71.91	69.88	37.23
AQLoRA-q	72.20	69.55	38.57
rsLoRA	66.58	37.40	36.89

D The rsLoRA composition

Budget under the composition. The recommendation does not extend to the rsLoRA composition, and the evidence does not show it fails there either. Repeating the two cheap budgets with +rs gives 75.46 at ρ=.05 and 75.13 at ρ=.10, against 75.50 over five seeds at ρ=.20. Against the three-seed standard deviation of that arm (0.06) both read as clearly worse. At five seeds the standard deviation is 0.35 and both sit within one of the ρ=.20 mean, so the comparison separates nothing. One seed per budget against an arm this variable is not evidence. The composition stays at ρ=.20 for want of evidence, not because the cheaper budgets were shown to fail.

AQLoRA can be composed with rsLoRA scaling (Kalajdzievski, 2023). The effect is model-dependent, and on six seeds per arm it does not reach significance. On phi-1.5 AQLoRA+rs averages 75.49 ± 0.31 against the rsLoRA baseline's 75.16 ± 0.37, a margin of 0.32 with p=.14 and overlapping ranges (Table 12).

How that number moved is the more useful result. At three seeds the margin was 0.29 with p=.005 and no overlap between runs. At six seeds it is 0.32 with p=.14. The effect size barely changed. Three seeds put both standard deviations near 0.06; six put them near 0.35. The significance was in an underestimate of the spread, not in the effect, and nothing in the three-seed version signalled it.

We therefore do not claim a win here. The margin is positive on the seeds we have, it is of the same size as the spread of the arms producing it, and its scope was already narrow: the same scaling collapses on Llama-3.2-1B (below) and at r=32 on phi-1.5 (§C). A practitioner should read +rs as a composition worth trying with a tuned learning rate, not as a property of AQLoRA. On Llama-3.2-1B it hurts badly. rsLoRA scaling divides the LoRA scaling factor by $\sqrt{r}$ instead of by r, which enlarges the update 2.8× at r=8. At our learning rate of 3×10^{-4} that overshoots on Llama: every rsLoRA-scaled run (AQLoRA+rs, rsLoRA, AQLoRA-s+rs) fell to 65 to 67, while the same methods without the scaling trained normally (the 71–73 range of Table 9). We did not tune the learning rate per method, so these runs show the failure, not whether a smaller rate would avoid it. The rank sweep (§C) reaches the same failure from the other direction: on phi-1.5, where +rs helps at r=8, raising the rank to 32 enlarges the same factor to 5.7× and drives the rsLoRA baseline to chance. We therefore do not scale by default and treat +rs as an optional add-on that needs its own learning rate per model and per rank.

Table 12: Seeds 42 to 47 on phi-1.5, 8-task average (%).

Method	42	43	44	45	46	47	mean ± std
QLoRA	74.74	74.71	74.80	74.30	74.62	74.97	74.69 ± 0.22
LoftQ[†]	75.13	74.65	74.73	–	–	–	74.84 ± 0.26
AQLoRA-q	74.88	75.11	74.35	74.78	74.75	74.82	74.78 ± 0.25
rsLoRA	75.40	75.29	75.27	74.44	75.11	75.46	75.16 ± 0.37
AQLoRA+rs	75.56	75.61	75.67	75.75	74.89	75.42	**75.49 ± 0.31**

[†]LoftQ was run on seeds 42 to 44 only.

Table 13: Edge profile (phi-1.5, batch 1, 256 tokens, rank 4). Device rows are projections.

Device	DRAM BW	proj. tok/s	peak GiB	1k-example epoch
V100 (**measured**), AQLoRA-edge	900 GB/s	232	1.44	6.8 min
V100 (**measured**), QLoRA-edge	900 GB/s	187	1.24	8.4 min
bandwidth-bound projection for AQLoRA-edge (upper bound; on-device kernel maturity will lower it):				
Jetson AGX Orin (LPDDR5)	205 GB/s	∼53	1.44	∼0.5 h
iPhone 16 Pro (LPDDR5X)	85 GB/s	∼22	1.44	∼1.2 h
OnePlus 15 / SD 8 Elite (LPDDR5X)	80 GB/s	∼21	1.44	∼1.3 h

E Edge feasibility

This section is a feasibility estimate, not a result: the only measurements in it are on the V100, and every device figure is a bandwidth-scaled projection. It bounds what the recipe would need on hardware not available to us, and it should not be read as more than that. A phone-sized run (batch 1, 256 tokens, rank 4) fine-tunes phi-1.5 in 1.44 GiB on a V100 at 232 tokens/s. That is 24% above QLoRA's throughput, at 1.44 against QLoRA's smaller 1.24 GiB, and well within the 8 to 24 GB of a 2026 flagship phone (Table 13). At batch 1 the work is bound by memory bandwidth, not compute: the run reaches 1.29 TFLOPS, 1% of the card's 125 TFLOPS fp16 peak, because it reads the weights once per token. So the right thing to scale by, going to a phone, is DRAM bandwidth. The V100 has 900 GB/s and the two phones in Table 13 have 80 and 85 GB/s, a factor of 10.6 to 11.3, which puts a phone at 1.2 to 1.3 h for a 1000-example epoch. QVAC reports 1 h 40 min per epoch for a larger 1.7B model on an Adreno 830 (Tether AI, 2025), an order-of-magnitude check on the projection rather than a match (their epoch is over a different dataset, so the two are not directly comparable). These are upper bounds; better mobile kernels will lower them. The on-device system, with per-tensor GGUF precision matching our precision map and a llama.cpp or Metal training path (Tether AI, 2025; Song & Tang, 2025), is future work.

F Timing forensics: two ways a sweep fails

One detail invalidated four sweeps, and it is easy to repeat. The first protocol timed every model over the same number of steps. A fixed step count is not a fixed measurement: Llama-3.2-1B is the fastest model here at 0.37 steps/s, so 60 measured steps took 164 s, while the same 60 steps on Qwen2.5-7B took 297 s. The same absolute jitter is then a much larger fraction of the shorter measurement, and the noise floor followed the window length rather than anything about the machine: 0.1% at 297 s, 1.8% at 217 s, and 3.1% at 164 s. Four Llama-3.2-1B sweeps failed the threshold and were discarded. The cause was the short window, not contention: Holding the *duration* fixed instead, at 200 measured steps for a 524 s window, brought the floor to 1.78% on the first attempt, and the resulting numbers agree with the discarded ones to within a third of a point. We therefore set the step count per model to target a comparable measurement window, and we recommend that anyone timing models of different speeds do the same.

The window is not the only thing that widens a floor. A later sweep on Llama-3.2-3B ran a 1006 s window on a card that read 20 MiB before all 42 runs, and still came out at 5.15%. Its per-repeat floors were 20.7, 11.3, 2.3, 1.7, 2.8, 8.1 and 0.3 percent, and every method slowed together partway through, which is what

Table 14: Precision budget and base WikiText-2 perplexity, phi-1.5.

Quantizer	fp16	ρ=0 (uniform)	ρ=.10	ρ=.20
NF4 (+DQ)	18.52	18.97	18.97	**18.94**
fp4 (no DQ)	18.52	19.43	19.38	**19.37**
fp16 layers	all	0	14	29
base (GiB)	2.64	0.95	1.05	1.18

interference looks like and what a method difference does not. We discarded it, and a later sweep on a quieter node reached 2.24% and is the one reported in Table 3. The two failure modes are separate: too short a window makes the noise unmeasurable, and a busy machine makes it large. Only the first is under the experimenter's control, which is why the floor has to be measured every time rather than assumed from an earlier session.

G Precision budget and base perplexity

A larger protection budget ρ lowers quantization damage on the base model. Table 14 gives the WikiText-2 perplexity of the quantized model with no adapters, so it measures the quantization alone, as ρ grows, under NF4 and under the harsher fp4 without double quantization. Keeping 20% of the layers in fp16 recovers a small, consistent part of the gap to fp16, and the effect is larger under fp4, where the budget matters most. Section 7.2 shows that this recovery depends on the budget, not on which layers are selected. The table also prices the budget: ρ=.20 moves 29 of 144 layers to fp16 and grows the base from 0.95 to 1.18 GiB, the 0.22 GiB premium that shows up again as peak memory in every phi-1.5 table.

H A residual protection tier

Protection can also be made cheaper instead of smaller. Storing a protected layer as a second NF4 stage over its own residual, rather than in fp16, removes 110 times the protected-layer reconstruction error for 8.25 bits per parameter against fp16's 16; halving the block size, which also adds scales, removes only 1.1 times. It works downstream and is still not worth using: 74.97 at 3.91 GiB peak, where plain fp16 at ρ=.05 scores the same at 3.89, and it gives up 3.6% step time because a protected layer then dequantizes twice instead of not at all. Given a memory budget, protecting fewer layers exactly beats protecting more layers approximately.

I Baseline fidelity notes

LoRA and DoRA. LoftQ runs on phi-1.5, Llama-3.2-1B and Phi-3-mini, and fails on Llama-3.2-3B and Qwen2.5-7B, where PEFT's initialisation cannot resolve the checkpoint on disk under our offline cache and raises before any weights are touched. We report it as an environment failure rather than a property of the method, because it does not follow model size and does not follow checkpoint layout: Phi-3-mini is sharded exactly as the two failing models are and initialises normally, so a single-file checkpoint requirement does not explain it either. A fix would need network access to re-fetch the affected snapshots, which the cluster does not have. LoftQ is the closest competitor to AQLoRA in this study and is level with it wherever both run, so its absence at two scales is a gap in our evidence and we do not read it in our own favour.

Fidelity notes (each baseline's settings were checked against its official repository). All methods share the protocol above, which differs from some originals. QLoRA's paper uses r=64 with scaling factor 16 (a scaling-to-rank ratio of 0.25 against our 2.0), bf16 compute, a constant learning rate, and a paged 8-bit optimizer. The optimizer only affects memory, so we use a fused AdamW for every method. Our *DoRA* rows sit on the quantized base, which is the accepted **QDoRA** form (Liu et al., 2024) and is not comparable to NVlabs' fp16 numbers. LoftQ uses PEFT's one-shot initialization; its headline gains are at 2–3 bit, so parity with QLoRA at NF4 is the expected result.

Table 15: Which layers to adapt, phi-1.5 (8-task avg, %).

Placement	budget α	mean block	seed 42, +rs	seed 43, no rs
near-zero frac (ours)	0.5	12.5	73.6	72.4
weight-norm	0.5	12.5	74.2	73.0
depth-matched random	0.5	12.2, 12.6	74.3	73.3
random	0.5	10.9, 11.3	74.6	73.5
near-zero frac (ours)	0.7	13.4	74.8	73.7
weight-norm	0.7	13.4	75.1	73.8
random	0.7	10.8, 11.7	75.1	74.2
depth (top blocks)	0.7	15.2	71.2	n/a
depth (top blocks)	0.85	13.3	74.6	n/a
near-zero frac (ours)	1.0	11.5	75.6	n/a

ρ=.2. The two columns differ in seed and in rsLoRA scaling, so compare within a column. "Mean block" is the average block index of the adapted layers, 0 to 23; random draws a new set per seed, so it has two. A second depth-matched draw without scaling, at seed 42, scores 73.5. Without scaling, depth at α=.85 is the shipped AQLoRA-s (73.2) and α=1.0 is AQLoRA-q (74.9 at seed 42).

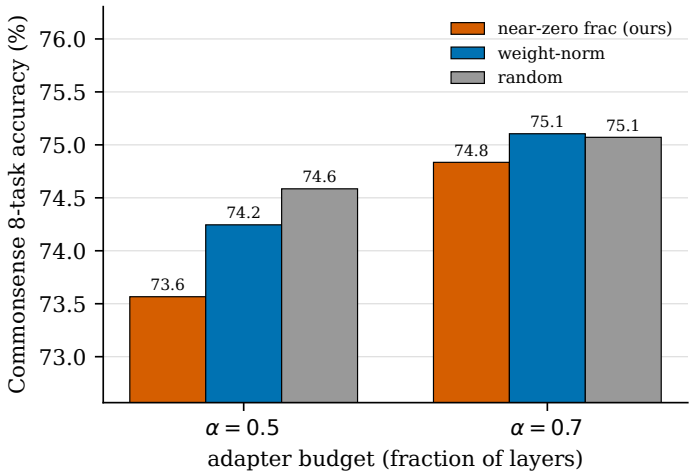


Figure 4: Layer-selection rules against random, two budgets (phi-1.5, seed 42 with rsLoRA scaling; Table 15 has both seeds).

J Adapter placement: figure and full table

K Lever ladder

L Full lever data

Table 17 lists all sixteen configurations (phi-1.5, 220 steps, one V100, one sequential pass): the fourteen lever rungs and two edge profiles that are the source for Table 16.

M Protection budget sweep

Table 18 sweeps the protection budget on phi-1.5. The three budgets we timed, in one session, span 1.1%, well inside that session's 2.8% floor, so we cannot separate them on speed and do not claim a trend.

This result does not carry to every setting. The sweep above is phi-1.5 with the quality setting. Under the rsLoRA composition the budgets stop being interchangeable, in a way one seed each cannot resolve. On Phi-3-mini the cheap budget is also worse: ρ=.05 scores 86.78 against 87.04 at ρ=.20 and 86.99 for QLoRA,

Table 16: Speed levers on phi-1.5 (220 steps, V100). GC is gradient checkpointing.

	Configuration	steps/s	tok/s	pad waste	peak (GiB)
universal	QLoRA, eager attention	0.228	1,046	39.8%	3.79
	+ SDPA	0.231	1,060	39.8%	3.79
	+ dropout 0	0.216†	993	39.8%	3.79
	+ length-grouped batching	0.245	1,124	**3.7%**	3.79
	+ fused AdamW (= QLoRA best)	0.247	1,134	3.7%	3.83
	+ no GC	0.399	1,833	3.7%	18.45
AQLoRA	sparsity α=1.0 (+rs, levers)	0.260	1,196	3.7%	4.05
	depth α=.7, no truncation	0.287	1,319	3.7%	4.02
	depth α=.7 + truncation	**0.303**	**1,392**	3.7%	4.02
	depth α=.5 + truncation	0.338	1,554	3.7%	4.00
	depth α=.7 + trunc., no GC	**0.568**	**2,609**	3.7%	14.10
	fp16 LoRA (levers)	0.339	1,558	3.7%	5.18
	rsLoRA (levers)	0.242	1,112	3.7%	3.83

These are a single sequential pass, used to attribute the speedup to its parts; the measured comparisons are Tables 3 and 4. The ladder uses depth α=.7 for attribution; the shipped speed setting is α=.85. The AQLoRA rows include rsLoRA scaling, a scalar on the update that does not change step time. †This row came out slower than the one above, within single-pass run-to-run noise.

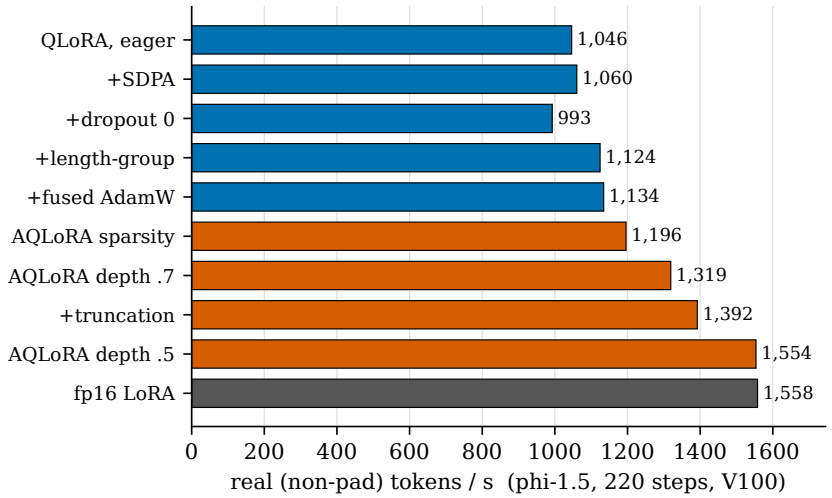


Figure 5: Throughput on phi-1.5 as levers are added. Blue: shared levers; orange: AQLoRA's own parts; grey: fp16 LoRA.

over three seeds with standard deviations near 0.15. The drop is small and inside its own spread, but it is in the wrong direction, so ρ=.05 is not a default. The narrower statement holds: the budget is worth sweeping per model, and on the model swept in full the cheap end lost nothing. The main tables report ρ=.20 because that is the budget the rest of the experiments used.

Appendix D covers the rsLoRA composition.

N Error ranking against matched random, per setting

Every row holds the protected parameter total equal to top-K's to within 0.1%, so the memory is matched and only the selection differs. Perplexity is base WikiText-2 with no adapters, which is deterministic given the layer set. The count K is the size of the top-K set; the matched draws vary in count because they are matched on parameters instead.

Table 17: Full lever data on phi-1.5 (measured): fourteen lever rungs and two edge profiles.

Configuration	steps/s	real tok/s	pad waste	peak GiB
QLoRA, eager attention	0.228	1,046	39.8%	3.79
+ SDPA	0.231	1,060	39.8%	3.79
+ dropout 0	0.216	993	39.8%	3.79
+ length-grouped batching	0.245	1,124	3.7%	3.79
+ fused AdamW (= QLoRA best)	0.247	1,134	3.7%	3.83
+ no grad-ckpt	0.399	1,833	3.7%	18.45
AQLoRA all-layer 1.0 (paged 8-bit opt)	0.238	1,093	39.8%	4.01
AQLoRA all-layer 1.0 + levers	0.260	1,196	3.7%	4.05
AQLoRA depth .7, no truncation	0.287	1,319	3.7%	4.02
AQLoRA depth .7 + truncation	0.303	1,392	3.7%	4.02
AQLoRA depth .5 + truncation	0.338	1,554	3.7%	4.00
AQLoRA depth .7 + trunc, no GC	0.568	2,609	3.7%	14.10
fp16 LoRA + levers	0.339	1,558	3.7%	5.18
rsLoRA + levers	0.242	1,112	3.7%	3.83
AQLoRA edge (bs1/seq256/r4)	0.154	232	0.0%	1.44
QLoRA edge (bs1/seq256/r4)	0.124	187	0.0%	1.24

Table 18: Protection budget swept end to end, phi-1.5.

ρ	fp16 layers	base GiB	peak GiB	8-task avg	step time
0 (QLoRA)	0	0.95	**3.83**	74.75 ± 0.05	ref
.05	7	1.01	3.89	74.97	+4.1%
.10	14	1.05	3.93	**75.08**	+4.8%
.15	22	1.14	4.01	74.48	
.20	29	1.18	4.05	74.78 ± 0.39	+5.2%

Rows ρ=0 and ρ=.20 are means of seeds 42 to 44; the other rows are seed 42. Step time is from one session, relative to QLoRA, with a 2.8% floor.

Table 19: Error ranking against parameter-matched random draws, per setting.

Model	ρ	K	top-K	bottom-K	matched random (min/med/max)	draws < top
phi-1.5	0.10	14	18.969	18.974	18.911 / 18.936 / 18.942	10/10
phi-1.5	0.20	29	18.935	18.918	18.871 / 18.906 / 18.941	8/10
phi-1.5	0.30	43	18.864	18.869	18.833 / 18.884 / 18.909	3/10
Llama-3.2-1B	0.10	11	9.492	9.476	9.463 / 9.490 / 9.507	6/10
Llama-3.2-1B	0.20	22	9.457	9.389	9.432 / 9.471 / 9.502	3/10
Llama-3.2-1B	0.30	34	9.412	9.268	9.388 / 9.439 / 9.457	2/10
Llama-3.2-3B	0.10	20	7.320	7.289	7.317 / 7.324 / 7.335	1/10
Llama-3.2-3B	0.20	39	7.308	7.249	7.241 / 7.310 / 7.320	4/10
Llama-3.2-3B	0.30	59	7.289	7.208	7.200 / 7.279 / 7.296	8/10
Phi-3-mini	0.10	13	5.647	5.677	5.605 / 5.657 / 5.665	3/10
Phi-3-mini	0.20	26	5.602	5.658	5.607 / 5.629 / 5.638	0/10
Phi-3-mini	0.30	38	5.538	5.626	5.544 / 5.592 / 5.625	0/10
Qwen2.5-7B	0.10	20	6.283	6.287	6.288 / 6.292 / 6.294	0/10
Qwen2.5-7B	0.20	39	6.271	6.277	6.264 / 6.284 / 6.290	2/10
Qwen2.5-7B	0.30	59	6.238	6.265	6.240 / 6.263 / 6.273	0/10

O Timing estimator and exclusion sensitivity

The headline speed numbers use one estimator (mean of the fastest half of the repeats) and one inclusion rule (duplicate-arm floor under 3%). Table 20 recomputes the pooled result under three other estimators, with the five excluded sessions forced back in, and on two subsets of the kept sessions: the nine with a duplicated

Table 20: Pooled speed vs. QLoRA under alternative timing choices.

Estimator	Sessions	n	AQLoRA-q	AQLoRA-s
fastest half	clean (paper)	13	$+4.8 \pm 2.0$	$+11.3 \pm 2.3$
fastest half	all	18	$+4.5 \pm 4.2$	$+10.0 \pm 7.8$
plain mean	clean	13	$+5.0 \pm 2.1$	$+11.2 \pm 3.4$
plain mean	all	18	$+5.2 \pm 4.6$	$+10.3 \pm 8.0$
median	clean	13	$+5.3 \pm 2.3$	$+11.6 \pm 4.4$
median	all	18	$+4.9 \pm 4.4$	$+10.3 \pm 8.6$
single fastest	clean	13	$+4.2 \pm 2.7$	$+10.8 \pm 2.1$
single fastest	all	18	$+4.7 \pm 4.2$	$+10.6 \pm 4.4$
fastest half	duplicate arm only	9	$+5.4 \pm 1.9$	$+11.7 \pm 2.3$
fastest half	empty-GPU check only	12	$+4.7 \pm 2.0$	$+11.1 \pm 2.4$

"All" adds the five floor-failing sessions, including the starved two-repeat Phi-3-mini session. The last two rows are subsets of the thirteen kept sessions.

QLoRA arm, and the twelve that waited for an empty timing GPU before each run. The point estimate moves by at most 1.3 points and never changes sign or ordering. What the exclusion rule buys is not the effect but the error bar: including sessions that fail their own floor more than triples the spread, from ± 2.3 to ± 7.8 on AQLoRA-s, without moving the mean beyond its deviation. Every number recomputes from the released `timing.json` files.