THE UNDETECTED DAMAGE OF QUANTIZATION ON RETRIEVAL AND HOW TO FIX IT

Anonymous authors
Paper under double-blind review

ABSTRACT

We show that a quantized model that keeps its classification accuracy still changes 14 to 46% of its top-1 retrieval results, and that aggregate ranking metrics reveal only part of this damage. We tie this failure to the gap between the two highest model scores and use that gap to decide when to trust a quantized answer and where additional precision should be spent. We show that the top-1 result is guaranteed to survive quantization only when this gap exceeds twice the largest rounding error. In classification, the scores are logits, and training compares the correct class against every other class, which encourages this gap. In retrieval, the scores are query-document similarities, and training compares each positive only against sampled negatives, so nothing separates the top-1 item from the second. This gap can be measured without labels. Before deployment, it predicts which models will break under quantization, and at deployment time it tells, per input, whether the quantized answer still matches the full-precision answer. Most classification inputs have a gap wide enough to trust the quantized answer, but few retrieval queries do. That gap motivates a different fix in each task. In retrieval, spending extra bit-width on the layers whose quantization moves the gap most recovers up to three-quarters of an extra bit's benefit for half its cost. In classification, routing the few low-gap inputs to full precision recovers most of the lost accuracy at a fraction of the cost.

1 INTRODUCTION

Quantization compresses a neural network by rounding its weights to a lower bit-width. Post-training quantization (PTQ) (Gholami et al., 2022) is often evaluated by how much accuracy the model loses relative to full precision. A 4-bit model that nearly matches its full-precision checkpoint in accuracy is declared a success, and for classification that intuition is right. But accuracy measures one thing only: whether the top-1 prediction changes, yet most deployed models are not classifiers. Retrieval, retrieval-augmented generation, recommendation, and reranking all depend on an *ordering*. For them, we find that a quantized model that preserves full-precision classification accuracy can still rank items very differently, including at the top-1. This problem goes undetected by metrics of quantization performance.

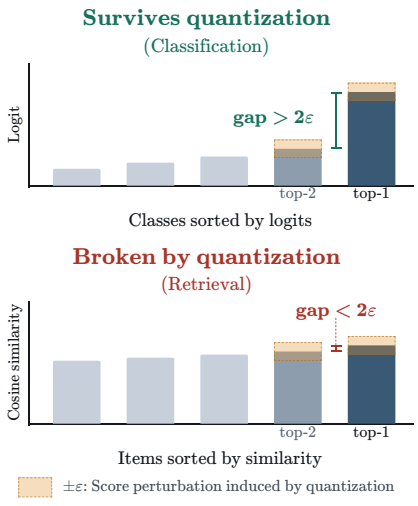


Figure 1: What quantization preserves and breaks. Every score moves by at most ε, so a top-1 result survives whenever it leads the runner-up by 2ε.

The effect is large and consistent. Across three backbones and two modalities, 4-bit round-to-nearest weight quantization (W4) changes the top-1 prediction of a fine-tuned checkpoint for only 3–6% of inputs when it is used as a classifier, but for 33–46% of queries when the same model is used for retrieval. The aggregate metrics understate this issue. On CLIP (Radford et al., 2021) text-to-image retrieval, Recall@1 falls by only 1.3% while about one in ten correctly answered queries loses its image. On Qwen3-Embedding-8B (Zhang et al., 2025), nDCG@10 drops by only 3.1% while one in six queries with a relevant top result loses it.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

The reason is which candidates each task family trains against. A ranking is stable under quantization only where the gap between two scores exceeds twice the worst-case rounding error, whether the scores are logits in classification or query-document similarities in retrieval. Classification trains against a fixed, complete set of classes, so every class is pushed below the correct one at every step, and the top-1 logit ends up far from the runner-up. Retrieval trains against a sampled subset of documents, so a positive is separated only from the negatives it was paired with, and nothing pushes the top document away from the second (Figure 1). To make this exact, let ε be the largest amount by which quantization perturbs any score. Since the top-1 can fall by ε while a challenger rises by ε, a candidate can overtake the top-1 only if it falls behind by less than 2ε. When no candidate is that close, the top-1 cannot change, on the other hand, when some are, the number of them predicts how likely the result changes. For most classification inputs, no candidate is that close, while for nearly every retrieval query, several are. This is why the same quantization leaves classification accuracy intact but breaks retrieval rankings. The same gap accepts or rejects individual answers at deployment time, without labels. We compare the quantized model's own top-1/top-2 gap against a single threshold, calibrated in advance on unlabeled data. The comparison tells us, for each input, whether the quantized answer is likely to match that of full precision model. Most classification inputs pass this check, but few retrieval queries do, given their typically small gaps. The damage is not an artifact of naive rounding, as quantizers built to minimize rounding error, such as GPTQ (Frantar et al., 2023), AWQ (Lin et al., 2024), HQQ (Badri & Shaji, 2023) and AdaRound (Nagel et al., 2020), still show the same failure.

Since nearly every retrieval query is at risk, the intervention must apply before deployment. In a mixed-precision setting, we allocate extra bit-width to the layers whose quantization most perturbs the gap between the top two items. In classification, only a small fraction of inputs is at risk, so the intervention can be adaptive with a first quantized forward pass estimating the gap, and one threshold on it routes the at-risk inputs to the full-precision model. Contributions are summarized below:

1. **Quantization breaks retrieval when accuracy holds.** When we hold a checkpoint fixed and change only the task, the top-1 retrieval result at W4 changes 5–23$\times$ more often than the classification answer (Tables 1 and 11). Standard retrieval metrics fail to detect this change. The damage is not tied to the quantizer, since stronger ones do not solve the issue.
2. **One label-free quantity explains and predicts it.** The gap between the top two scores, measured against the perturbation, explains why retrieval breaks where classification does not, predicts the damage across models, quantizers, and corpora, and tests individual inputs without labels.
3. **Two fixes from one quantity.** On retrieval, allocating bit-width by gap sensitivity recovers three-fifths to three-quarters of an extra bit's benefit for half its cost and outperforms every criterion under every quantizer we test. On classification, the same gap, thresholded once, routes the inputs at risk to full precision, recovering most of the lost accuracy at a fraction of the full-precision cost.

2 Related work

Post-training quantization. Post-training methods take a trained model and fit the quantization to a small unlabeled calibration set, choosing scales, rounding, or a rotation of the weights so that each quantized layer reproduces the outputs of the original one: GPTQ (Frantar et al., 2023), AWQ (Lin et al., 2024), SmoothQuant (Xiao et al., 2023), QuaRot and SpinQuant (Ashkboos et al., 2024; Liu et al., 2024), and for vision transformers PTQ4ViT, RepQ-ViT and AdaLog (Yuan et al., 2022; Li et al., 2023b; Wu et al., 2024), surveyed by Gholami et al. (2022).

Ranking under compression. On mixture-of-experts routing, rank-preservation losses keep a quantized router choosing the experts that full precision would choose. A separate analysis notes that a small perturbation from quantizing the KV cache can change which experts top-k routing selects, but it does not measure how often (Park et al., 2026; Gu, 2026). On classification, perturbation bounds have been stratified by the top-2 logit gap for the true-label margin (Kiselev, 2026). On ranking under uncertainty, when scores are intervals, two items can be compared only if their intervals do not overlap (Soliman & Ilyas, 2009). When every interval has the same width, the relation is a semiorder (Luce, 1956; Scott & Suppes, 1958; Fishburn, 1970), and the rankings a bounded perturbation can reach are its linear extensions (Szpilrajn, 1930). We restate the top-1 case in the notation the rest of the paper uses. What is new here is the measurement, not the condition. On embedding compression, the stored document vectors are quantized (Jégou et al., 2011; Guo et al.,

2020) or the query encoder is trained jointly with the index (Zhan et al., 2021), and reconstruction error is known to be the wrong objective for ranking (Guo et al., 2020). We differ in three ways. We quantize the encoder itself rather than the index, we read one model both as a retriever and as a classifier, and we test individual inputs at deployment time without labels.

Mixed-precision allocation. Per-layer bit-width assignment is established for classification (Dong et al., 2019; 2020; Yao et al., 2021), with ZeroQ replacing curvature by a KL divergence to the full-precision outputs (Cai et al., 2020) and mixed precision reaching the KV cache of long-context inference (Tao et al., 2025). All target output errors. None measures what a layer's quantization does to the ranking in retrieval.

Selective prediction and routing. Selective prediction sends inputs a model is unsure about to a stronger one, scored by maximum softmax probability (Hendrycks & Gimpel, 2017; Geifman & El-Yaniv, 2017) or by the margin between the best and second-best class, first used as an uncertainty score for active learning (Joshi et al., 2009). Cascades, early exits, and speculative decoding with quantized drafts rely on the same mechanism (Teerapittayanon & McDanel, 2016; Huang et al., 2018; Kolawole et al., 2024; Georganas et al., 2025; Zhao et al., 2025). Compression changes predictions while accuracy holds (Hooker et al., 2019; 2020; Dutta et al., 2024), and among single scores the margin is the best predictor of which predictions change (Hu et al., 2022).

3 When a ranking can change

The rounding error introduced by quantizing the weights propagates through the forward pass and displaces every score by a bounded amount. Although we measure the top-1 throughout, we analyze entire rankings because a retrieval top-1 is the first position of a ranking over thousands of candidates, and what displaces it is often a candidate from deep in the list rather than the runner-up. Formally, let $z \in \mathbb{R}^N$ be the *score vector* over N candidates, with z_i the score of candidate i. Throughout, a *score* is whatever quantity the system ranks by, a class logit in classification and a query-candidate similarity in retrieval. The analysis below applies to both cases. A *permutation* $\sigma = (c_1, \dots, c_N)$ lists the candidates, so that c_k is the candidate in k-th position. The *ranking* of z, denoted $\pi(z)$, is the permutation that sorts z, *i.e.* the one for which $z_{c_1} > \cdots > z_{c_N}$, and it is undefined when z has ties. Given $\varepsilon > 0$, a *perturbation* is any $\delta \in \mathbb{R}^N$ with $\|\delta\|_\infty \le \varepsilon$.

A permutation σ is *ε-reachable* from z if some perturbation δ has $\pi(z+\delta) = \sigma$. We write $\mathcal{R}_\varepsilon(z)$ for the set of reachable permutations. It always contains $\pi(z)$, via $\delta = 0$. We model quantization as replacing z with $z + \delta$ at the output, where ε depends on the bit-width.

Define the k-th *gap* as $g_k(z) = z_{c_1} - z_{c_k}$, the difference between the first score and the k-th, so that $g_1 \equiv 0$ and g_k is non-decreasing in k. Since a perturbation can change one score by at most $\pm\varepsilon$, the k-th candidate can overtake the winner only if $g_k(z) < 2\varepsilon$. We therefore call $\mathcal{C}_\varepsilon(z) = \{\, k \ge 2 : g_k(z) < 2\varepsilon \,\}$ the *contender set*, *i.e.* the positions that a perturbation could possibly push past the top-1, and $\mathcal{C}_\varepsilon(z) = \emptyset$ means the top-1 has no contender and cannot be surpassed.

Given an input x, we write $z(x)$, $g_k(x)$ and $\mathcal{C}_\varepsilon(x)$ for the scores, gaps and contenders it determines. These are obtained by running both the full-precision (FP) model ($\cdot^{\mathrm{FP}}$) and its quantized counterpart ($\cdot^{\mathrm{PTQ}}$) on x and measuring, not assuming, the perturbation $\delta = z^{\mathrm{PTQ}} - z^{\mathrm{FP}}$, from which we take $\varepsilon = \|\delta\|_\infty$. We call $g_2(x)/2\varepsilon$ the *separation ratio*. It is at least one exactly when $\mathcal{C}_\varepsilon(x) = \emptyset$, so we call one the *stability threshold*. Across inputs, a median separation ratio above one means the typical input has a stable top-1. Every diagnostic below computes it from the FP gap and the perturbation.

Lemma 1 (Top-1 and top-k stability)**.** *Write $\pi(z) = (c_1, \dots, c_N)$. For each k from 1 to $N-1$, every permutation $\sigma \in \mathcal{R}_\varepsilon(z)$ has the same first k elements as $\pi(z)$, as a set, iff $z_{c_k} - z_{c_{k+1}} \ge 2\varepsilon$. For $k = 1$ this says the top-1 survives every perturbation iff $g_2(z) \ge 2\varepsilon$, that is, iff $\mathcal{C}_\varepsilon(z) = \emptyset$.*

Each score moves by at most ε, so a candidate can overtake another only if that other started at most 2ε ahead of it. When the gap is smaller, the proof of Lemma 1 in Appendix C constructs a perturbation that swaps the two. Both conditions are known. They restate, in the notation of this paper, the classical case in which each score is known only up to an interval of fixed width (Luce, 1956; Scott & Suppes, 1958; Fishburn, 1970; Szpilrajn, 1930). Appendix B states the general characterization of reachable rankings (Proposition 1) and a closed-form flip probability under a uniform perturbation model (Proposition 2). The top-k condition would be the natural guarantee for ranked retrieval, but

the gaps below the top-1 are an order of magnitude too small for it to be usable in practice (Table 18), so everything we propose rests on the top-1 condition. Being within 2ε is not transitive, so a chain of near-ties lets a candidate far below the top-1 overtake it, which is where a large share of retrieval flips come from. The first gap is still enough to diagnose the risk.

Corollary 1 (Contender locality)**.** *Let a be the top-1 of z. Candidate j overtakes a under perturbation δ iff $\delta_j - \delta_a > z_a - z_j$, so whether the top-1 survives depends on δ only through the differences $\{\delta_j - \delta_a\}$ and not on $\|\delta\|_\infty$ over all N coordinates. Consequently, writing a' for the top-1 of $z + \delta$ and $S = \max_j(\delta_{a'} - \delta_j)$, a change of top-1 implies $g_2(z + \delta) < S$, and $S \le 2\|\delta\|_\infty$.*

Only differences matter: a perturbation that moves every score by the same amount leaves every ranking untouched, however large the perturbation happens to be. The global bound 2ε ignores this since it also counts a shift that the winner and a challenger share, which cancels in their difference. That is why the stability check below calibrates S rather than 2ε.

From condition to stability check. The top-1 condition of Lemma 1, $g_2(z^{\mathrm{FP}}) \ge 2\varepsilon$, is stated with the full-precision scores and the size of the perturbation, and at deployment time only the quantized model runs, so neither is available. The condition can, however, be rewritten in terms of the quantized scores alone. Write c_1' for the candidate the quantized model ranks first, $g_2(z^{\mathrm{PTQ}}(x))$ for its top-1/top-2 gap, and $S(x) = \max_j(\delta_{c_1'} - \delta_j)$ for the largest amount by which δ moves a score difference involving c_1'. Whenever the top-1 changes, $g_2(z^{\mathrm{PTQ}}(x)) < S(x)$, and $S(x)$ never exceeds 2ε (Corollary 1). So if the quantized gap is at least $S(x)$, the top-1 did not change, and unlike ε, S can be calibrated once and stored. On n unlabeled calibration inputs, run both models, compute S on each, and let $\hat{\tau}$ be the $\lceil (n+1)(1-\alpha) \rceil$-th largest value, the split-conformal quantile. At deployment time, accept a test input, i.e., trust its quantized top-1 as the full-precision answer, iff $g_2(z^{\mathrm{PTQ}}(x)) \ge \hat{\tau}$. When the calibration and test inputs come from the same distribution, split-conformal calibration (Vovk et al., 2005; Lei et al., 2018) gives

$$\mathbb{P}\big(\text{accepted and top-1 changed}\big) \;\le\; \alpha,$$

with no labels at test time. Calibrating 2ε instead of S is also valid but accepts far fewer inputs, and no retrieval query at all (Appendix A.6). The calibration pass costs nothing extra, since a quantization pipeline already runs the FP model on unlabeled data to calibrate quantizers such as GPTQ.

4 Setup

Models. We study three fine-tuned backbones across two modalities: ViT-B/16 and ViT-L/16 (Dosovitskiy et al., 2021) for images, and Qwen3-Embedding-0.6B (Zhang et al., 2025) for text, giving 53 fine-tuned checkpoints in total. Later sections add five more text embedders, three encoder-based (Xiao et al., 2024; Wang et al., 2022; Li et al., 2023a) and Qwen3-Embedding-4B and 8B, together with CLIP (Radford et al., 2021), a cross-encoder reranker (Bajaj et al., 2018), and a Qwen3-8B generator for the retrieval-augmented pipeline. Quantization is done on weights at 4 and 3 bits (W4 and W3), with activations kept at full precision except in a specific experiment. Every linear layer of every model is quantized. Appendix A.2 details fine-tuning and the remaining settings.

Tasks. Classification and retrieval each run in both modalities, and a fifth setting is a deployed cross-modal system. In image classification, each ViT is fine-tuned separately on 21 image classification tasks. In text classification, Qwen3-Emb-0.6B is fine-tuned on 11 tasks from MTEB (Muennighoff et al., 2023). In image retrieval, the same fine-tuned ViTs rank each task's test split by similarity to a query image drawn from that split. This compares classification and retrieval on the same model. In text retrieval, the Qwen3 embedders rank four standard BEIR corpora (Thakur et al., 2021), NFCorpus (Boteva et al., 2016), SciFact (Wadden et al., 2020), SCIDOCS (Cohan et al., 2020) and FiQA (Maia et al., 2018). In text-to-image retrieval, CLIP ranks 4,000 Flickr30k images (Young et al., 2014) by similarity to a caption, with no fine-tuning. Section 5.3 adds two systems built on retrieval: a cross-encoder reranker on the same BEIR corpora and a retrieval-augmented generation pipeline on Natural Questions (Kwiatkowski et al., 2019).

Evaluation. We report the top-1 result throughout (the predicted class in classification, the first-ranked item in retrieval) and measure how often it differs from the FP model. This needs no relevance labels, since it compares the two models rather than either one against ground truth. Six procedures need calibration data, and none takes labels. The allocation criteria use 128 calibration queries by

default, and the data-dependent quantizers use 512 documents. Calibration and evaluation are disjoint throughout, with one exception described in the calibration protocol in Appendix A.2.

5 What survives quantization

The same fine-tuned model, under the same quantizer, at the same bit-width, changes its top-1 answer for a few percent of inputs when asked to classify and for over forty percent of queries when asked to retrieve (Table 1). The Qwen3 rows show the same asymmetry in another model family, first with a fine-tuned classifier against the pretrained embedder, then with one checkpoint read both ways. The retrieval top-1 changes 4.8–14.3× more often than the classification top-1. The perturbation is similar in both cases, and what changes is the gap it has to cross. Classification trains against every class and retrieval against a sampled subset, so a positive is separated only from the negatives it was paired with, and the second place is a comparison training never made.

Table 1: Same backbone, two tasks. ViT rows use one checkpoint for both tasks, the first Qwen3 row uses a fine-tuned classifier against the pretrained embedder, the last row uses one checkpoint, classified with a linear probe and retrieved by cosine self-retrieval. Top-1 change rate at W4 group_128 under RTN, mean ± standard error over n tasks.

Backbone	Classification	Retrieval	Ratio
ViT-B/16	6.2% ±1.4 (n=21)	45.5% ±2.5 (n=5)	**7.3**×
ViT-L/16	2.9% ±0.7 (n=21)	42.1% ±1.5 (n=5)	**14.3**×
Qwen3-Emb-0.6B	4.3% ±1.1 (n=11)	33.4% ±4.1 (n=4)	**7.8**×
Qwen3-Emb-0.6B, one checkpoint	8.6% ±3.1 (n=3)	41.2% ±4.9 (n=3)	**4.8**×

Most changed results are worse. Judged against the BEIR relevance labels over six embedders, four corpora and six quantizers, 66% of the flips whose FP top-1 was a relevant document replace it with one that is not, and 77% do at three bits (Table 20). A changed top-1 has two outcomes: either the new item is still relevant or it is not. So a 34% change rate does not by itself imply a loss of quality. On CLIP, where every query has one correct image, 26–30% of all changed results lose it and 13–23% gain it (Appendix A.7, *What a changed result costs*). The flips that keep a relevant item on top are not harmless either. A retriever is one component of a larger system, so a different document yields a different generated answer (§5.3), and two servers running the same model at different precisions answer the same query differently. The standard metric reveals only a fraction of this damage, harmful or not: nDCG@10 averages over ten positions, and the item that lost first place usually stays inside that window. Across three model families at W4, nDCG@10 falls by 2.7–7.6% while the relevant document leaves the top-1 for 16–22% of the queries that had one, and it leaves the top-10 just as often (Table 2). Accuracy also tells nothing about the rest of the ranking. On classifiers whose top-1 accuracy is preserved, the exact top-5 set survives for only 28-47% of inputs (Table 18).

Table 2: A mean quality metric understates per-query damage. *NDCG loss* is the relative change in mean nDCG@10. *Gold lost* is the fraction of queries for which a document that was both relevant and in the FP top-k falls out of the PTQ top-k. Round-to-nearest, group_128, mean over four BEIR corpora.

			Gold document lost from top-k		
Model	Bits	NDCG loss	k=1	k=5	k=10
Qwen3-Emb-0.6B	W4	−7.5%	22.2%	23.7%	23.3%
Qwen3-Emb-8B	W4	−3.1%	15.6%	16.2%	15.3%
BGE-large	W4	−2.7%	16.0%	17.8%	16.3%
E5-large-v2	W4	−7.6%	18.6%	21.6%	23.7%
Qwen3-Emb-0.6B	W3	−68.9%	79.6%	80.8%	78.4%

5.1 Nothing else explains the asymmetry

Nothing but the tasks explain why retrieval breaks more than classification. The asymmetry survives when everything else is held fixed. The last row of Table 1 embeds one dataset once with one encoder and reads the same vectors both ways, through a frozen linear probe and by cosine nearest neighbor, and at W4 retrieval still changes 3.3–6.8× more of its top-1 answers on each dataset (Table 7). Accuracy loss does not predict it. Within a quantization setting, the accuracy a configuration loses correlates with the fraction of top-5 sets it changes at only +0.327, and in the deployed W4-group_128 setting the correlation is undetectable on either vision backbone (Table 19). The number of candidates does not predict it either. Growing the corpus by up to 58× leaves the text change rate within three points and raises the vision rates sharply, and in both cases the top-1 change follows the separation ratio, not the candidate count (Table 21).

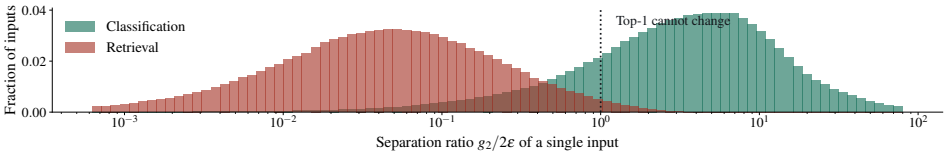


Figure 2: The two task families separate their winners by different amounts. The separation ratio of every input we measure, at W4-group_128, over both modalities. The dotted line is the stability threshold.

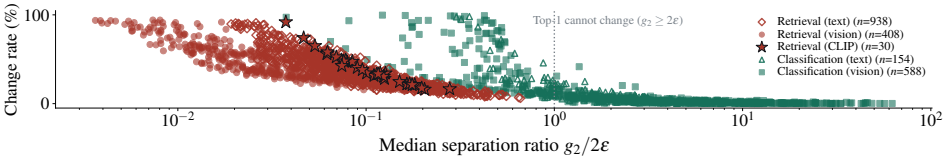


Figure 3: One quantity predicts the top-1 change rate in both task families. Each point is one setup (model, task, quantizer, bit-width), and a retrieval setup appears once per corpus size.

5.2 The contender set is the mechanism

The contender set $\mathcal{C}_\varepsilon$ is what separates the two tasks. It consists of the candidates within 2ε of the top-1, and at W4 it is empty for 74.3% of ViT-B/16 classification inputs but for none of the retrieval queries. The flips themselves are also informative. When classification changes its top-1, the new top-1 is the old runner-up 82.2% of the time, while in image retrieval it is only 39.7%, with the rest coming from deeper in the ranking, as a wider contender set predicts. The same holds on CLIP (Table 25). Over every input, 85.8% of classification inputs sit above the stability threshold against only 3.5% of retrieval queries, with median separation ratios of 5.73 and 0.081 (Figure 2). Among inputs with $|\mathcal{C}_\varepsilon| > 1$, the change rate rises with $|\mathcal{C}_\varepsilon|$ in both tasks, see Table 24.

The median separation $g_2/2\varepsilon$ tracks the change rate at Spearman -0.88 over 2,118 configurations, every one we measure (Figure 3). Within each family, the correlation is -0.86 (classification, text), -0.92 (classification, vision), -0.93 (retrieval, text), and -0.90 (retrieval, vision). Candidate count reaches only $+0.02$ on text and $+0.28$ on vision. The relationship survives intervention, which correlation alone cannot show. We fit a straight line of change rate against $\log_{10}$ of the median separation ratio by least squares. Fitted on the RTN image-retrieval points alone at W4 ($n = 34$, $R = -0.860$), the line predicts the same 34 points under each of the five other quantizers, which it never saw, at R^2 between 0.60 and 0.74. Changing the quantizer moves both quantities but keeps their relationship. The ratio divides by ε, and ε is also what causes the flips, so the correlation could seem trivial. It is not: on a single input, a changed top-1 always has $g_2 < 2\varepsilon$, so Figure 3 measures how well the gap predicts the damage, and neither quantity alone is as predictive. Alone, the median gap g_2 reaches Spearman -0.58 and the median perturbation ε just $+0.05$, against -0.88 for the ratio ($g_2/2\varepsilon$).

Table 3: The stability check on classification, text retrieval and CLIP. Coverage (mean ± s.d. over 20 calibration splits) and conditional violation rate at W4 group_128. The thresholds themselves are in Table 23.

Setting	α	Coverage	Violation
Classification (ViT-B/16)	0.10	85.2% ± 0.1	1.47%
	0.05	80.0% ± 0.1	0.72%
	0.01	67.7% ± 0.2	0.10%
Text retrieval (Qwen3)	0.10	11.3% ± 0.4	0.00%
	0.05	8.8% ± 0.3	0.00%
	0.01	4.8% ± 0.4	0.00%
CLIP ViT-B/32	0.10	1.5% ± 0.4	0.00%
	0.05	1.1% ± 0.3	0.00%
	0.01	0.5% ± 0.2	0.00%
CLIP ViT-L/14	0.10	15.7% ± 0.9	0.00%
	0.05	12.5% ± 0.9	0.00%
	0.01	7.8% ± 0.9	0.00%

The stability check of §3, which fixes the threshold in advance, is the deployable version that can fail, and Table 3 reports its violations. On text retrieval, it accepts about one query in nine, since the separation there is two orders of magnitude smaller, and on CLIP the outcome is similar (Table 3). The violation rate stays within α in every calibration split. This was not guaranteed because $\hat{\tau}$ is fixed on calibration data, so the rate could have exceeded α on test data (Appendix A.6).

5.3 The same effect in other systems

CLIP text-to-image retrieval, a fourth architecture family (Radford et al., 2021; Young et al., 2014), shows the failure on a deployed system: captions as queries, images as the corpus. ViT-L/14 at W4

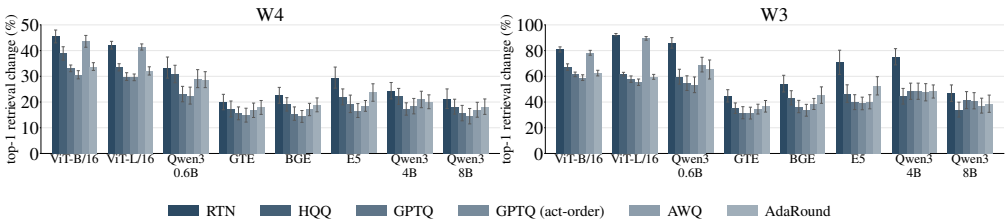


Figure 4: Stronger quantizers reduce ranking damage and do not remove it. Top-1 retrieval change rate at group_128, mean ± standard error over tasks. Numbers are in Table 10.

loses only 1.3% of its Recall@1, so the benchmark says quantization did little damage, while 24.6% of top-1 results change and one right answer in nine becomes wrong (Table 4). The metric stays flat because almost as many queries gain their correct image as lose it. The queries that keep it are the ones with the larger separation, and they change their top-1 four times less often.

Reranking, the second stage of a search system, breaks the same way. A cross-encoder takes each query and document together and rescores the shortlist the FP retriever returned, and we quantize only that cross-encoder (Bajaj et al., 2018) while keeping the shortlist fixed. The reranked top-1 changes for 5–10% of queries at W4 and 12–27% at W3, and the settings with more separation change fewer results (Table 9).

A *retrieval-augmented* pipeline returns a different answer to the user once its retriever is quantized, and at W4 its accuracy reports a twentieth of that change. We quantize only the query encoder, Qwen3-Emb-0.6B, and leave the generator, Qwen3-8B, at FP with greedy decoding, so every changed answer comes from retrieval, and we score 2,255 Natural Questions (Kwiatkowski et al., 2019) over 2.68M passages (results in Table 17). The retrieved passages carry the answer, since the same generator scores 15% with no passages against 43% with five. At W4, between 24 and 32% of the answers change under the six weight quantizers, while exact match by at most 1.6 points. More answers turn wrong than right, so the damage is real and the metric understates it. We report W4 only, because at W3 this encoder returns near-random passages and the pipeline does worse than the generator alone without even retrieval augmentation, which makes that setting unusable.

Encoder-based embedders follow the same rule, and separation orders models better than size does. With three encoder-based embedders (Xiao et al., 2024; Wang et al., 2022; Li et al., 2023a) beside the three Qwen3 models, the separation ratio predicts the change rate at Spearman -0.98 over the twelve model and bit-width configurations (Table 8), and GTE-large, with a twelfth of the parameters of Qwen3-Emb-4B, changes fewer top-1 results at both bit-widths. Thus, one should pick higher-separation models for retrieval, not larger ones.

5.4 Different quantization methods do not remove the asymmetry

The analysis so far is based on the RTN quantizer, and an objection could be that it is too weak. We therefore repeat the measurement with five others (Figure 4). GPTQ (Frantar et al., 2023) compensates rounding error against the activation covariance, with and without activation ordering, AWQ (Lin et al., 2024) rescales to protect outlier channels, HQQ (Badri & Shaji, 2023) fits scale and zero-point under a robust loss, and AdaRound (Nagel et al., 2020) learns the rounding direction of every weight. Settings are detailed in Appendix A.2, *Quantizers*. All five reduce the top-1 change rate, including on CLIP (Table 26), but none of them approaches the stability threshold on retrieval.

A finer quantization grid also does not solve the issue. We hold the bit-width fixed and change only how many weights share a scale, from one per output channel to one per 128 weights to one per 64. The rounding error falls, and the separation ratio $g_2/2\varepsilon$ rises under all six quantizers and in both task families, but the asymmetry widens (Table 16). Moving from the coarsest grid to the finest cuts the classification change rate by $3.7\times$, and the retrieval rate by $1.7\times$, so retrieval ends up changing 4.5 to $8.0\times$ more of its top-1 results than classification, against 1.7 to $5.9\times$ at the coarsest grid. A finer grid buys a real improvement in both families but moves neither across the threshold between them.

Quantizing the activations as well changes the size of the perturbation, but our results still hold. We add eight-bit activations, with one static range per tensor, on top of W4 and W8, under all six weight quantizers (Appendix A.10). The change rate still rises the same way, and retrieval's top-1 results still change seven to eight times more than classification. At W8A8, still no retrieval configuration

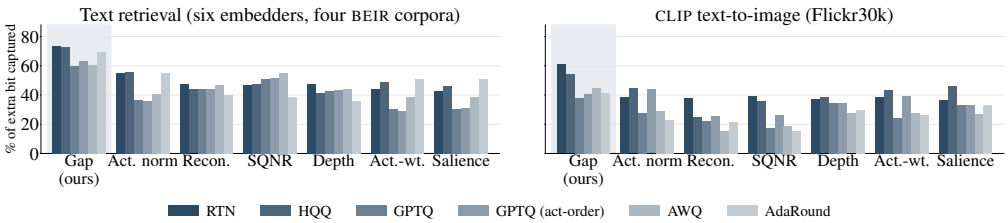


Figure 5: Gap sensitivity outperforms all other criteria we compare. Fraction of the W3→W4 reduction in top-1 change recovered at a 3.5-bit average, averaged across models and corpora.

clears the stability threshold while 91% of classification configurations do. A configuration's typical input does not flip when its median separation $g_2/2\varepsilon \geq 1$. At W3, GPTQ with and without activation ordering, HQQ and AdaRound push classification over the stability threshold, RTN and AWQ do not, and all six leave retrieval far below it. GPTQ on ViT-L/16 restores accuracy from 55.5% to 88.6% and cuts the change rate eleven-fold, while retrieval on the same model still changes most of its top-1 results (Table 11). Embedding retrieval sits one to two orders of magnitude below, too far for any quantizer to push it over.

6 Two fixes from one quantity

The gap suggests two fixes, and one number decides between them: the *at-risk fraction*, the share of inputs whose top-1/top-2 gap is below 2ε. At W4, it is 94–100% of retrieval queries but only 9–26% of classification inputs. Retrieval has nearly all inputs at risk, so in a mixed-precision setup, we allocate more bit-width to the layers whose quantization moves the gap most. Classification has only a minority at risk, and that minority can be detected without labels. So the fix is per input, and we route those inputs to the FP model and leave the rest on the quantized one. Both fixes need the same quantity, measured before deployment on unlabeled data.

6.1 Retrieval: allocating bits by gap sensitivity

Weight-only PTQ methods decide *how* to round and not *where* to spend precision. Bit-width is set once globally for all layers, and the goal is to minimize reconstruction error. Where to spend it is a separate choice, and the two combine freely: an allocation criterion decides which layers get more bits, and any quantizer then rounds them, so a criterion plugs into every quantizer. We propose an allocation criterion and compare it against baselines, under each quantizer fixed. Fix the mixed-precision *budget* at 3.5 bits, the average number of bits per weight the deployment allows. The question is, which half of the layers receive four bits and which receive three?

Method. Run the FP model on the $n = 128$ calibration queries against the corpus. For each query, this gives the top-1/top-2 gap $g_2(x)$ and the identities of the two documents involved. Then, for each linear layer ℓ, quantize that layer alone, with plain round-to-nearest (RTN) at 3 bits and group size 128, leave every other layer at FP, rerun the calibration queries, and compute the score difference between the *same two documents*, $g_2^{(\ell)}(x)$. The sensitivity of the layer is how much this one-layer quantization moves the gap on a typical query,

$$s_\ell \;=\; \mathrm{median}_x \big|\, g_2(x) - g_2^{(\ell)}(x) \,\big|.$$

Layers are ranked by s_ℓ divided by their parameter count, since bits are spent per weight, and the higher bit-width goes to the most sensitive layers until the budget is spent. The whole measurement is $L + 1$ forward passes over 128 queries for a model with L linear layers, needs no labels, and trains nothing. The sensitivities are measured once, with RTN, and the resulting assignment is then applied with whatever quantizer the deployment uses, so the method generalizes to unseen quantizers.

Baselines and metric. We compare gap sensitivity against allocating by reconstruction error, the objective of existing allocators, and five other criteria, defined in Appendix A.8. We report the fraction of the W3 → W4 benefit that an allocation captures while spending only half the extra bit at 3.5 bits, computed as $\sum$ gain/ $\sum$ headroom (the *capture ratio*).

Gap sensitivity leads. Gap allocation recovers 60–73% of the extra bit's benefit and is first among all seven forward-only criteria under every quantizer (Figure 5, Table 5). Its lead over the next best criterion ranges from 5.5 points (AWQ) to 18.4 (RTN), and it beats each competitor in at least 16 of the 24 configurations. In absolute terms, at the 3.5-bit budget under RTN, it changes 33.3% of top-1 results against 43.1% for reconstruction error and loses 13.8% of correct top-1 documents against 19.9% (Table 12), a mean paired gain of +26.2 capture points over the 24 configurations. Reconstruction error, the objective of existing allocators, captures only 40–48%. The reason is visible at the layer level: on the text embedders, reconstruction error is nearly uncorrelated with gap sensitivity across layers (mean Spearman -0.04, Table 14), and on CLIP, where the correlation reaches $+0.34$, it recovers more but still less than gap sensitivity.

Retrieval-augmented generation. Gap allocation also changes fewer answers under every quantizer. We run the same retrieval-augmented pipeline on Natural Questions at a 3.5-bit budget, change only the allocation, and repeat it for all six quantizers (Table 13). We compare gap allocation against reconstruction error, the most common objective that existing allocators optimize. On average, gap sensitivity recovers more of the extra bit under all quantizers by 18pp and leaves 62 fewer answers changed. Hence, gap allocation yields higher agreement between the quantized and FP models.

Robustness. The lead holds under every setting we vary. Across the six models, per-model capture stays between 68.3% and 84.5% (Table 14). It is unchanged at the other two budgets we test, 3.25 and 3.75 bits per weight, across calibration seeds, and at every calibration size from 8 to 256 queries, on text and on CLIP (Appendix A.9). On CLIP itself, over both encoders and all seven settings of budget, seed and quantizer, gap-driven allocation captures 56.4% of the extra bit against 34.9% for reconstruction error, in every configuration (Table 6, full results in Appendix A.7). The lead also survives quantizing the activations: with eight-bit activations on top of weight quantization, gap sensitivity captures 67% of the extra bit against 37% for reconstruction error (Appendix A.10).

6.2 CLASSIFICATION: ROUTING INPUTS BY THE GAP

In classification, the same gap that gives the stability check also provides a fix. Routing low-gap inputs to the FP model recovers most of the lost accuracy at a fraction of the FP compute.

Method. Let $g_{\mathrm{q}}(x)$ be the top-1/top-2 gap of the *quantized* logits. It comes at no extra cost, from the PTQ forward pass the deployment already runs. By the top-1 condition of Lemma 1, $g_{\mathrm{q}}(x) > 2\varepsilon$ means no perturbation of size ε changes the prediction, so at every budget the flippable inputs are exactly those with $g_{\mathrm{q}}(x) < 2\varepsilon$: one threshold on the gap separates them. Route x to FP iff $g_{\mathrm{q}}(x) < \tau$, with τ set at the 25th percentile of g_{q} on an unlabeled validation slice: no training, just one tuned hyperparameter. Up to order-equivalence, the gap is the only score of the quantized logits whose threshold sets match the worst-case flippable sets at every budget (Proposition 3).

Results. Routing 25% of inputs recovers 85–93% of the accuracy lost to quantization on the three fine-tuned backbones, over the tasks whose FP $\rightarrow$PTQ accuracy gap is at least 0.5 points (Table 15). The cost is one quantized pass over everything plus one FP pass over the routed quarter, 47% of the FP cost at the $4.5\times$ speedup GPTQ reports (Frantar et al., 2023). Recovery rises smoothly with the fraction routed, from 62–76% at 10% to 97–99% at 50% (Table 15). At nearly every operating point, the gap recovers more than MSP or predictive entropy (Gal & Ghahramani, 2016).

7 CONCLUSION

We show that the damage caused by quantization is tied to the score gaps between the top two candidates (logits in classification, query-document similarities in retrieval). When that gap is large, as in classification, the quantized model preserves its accuracy on most inputs. When the gap is small, as in retrieval, the quantized model changes many top-1 results, and standard metrics understate that damage. The gap is computable label-free and predicts which models will break before quantization and which answers to trust after it. The gap also motivates a fix in each task. In retrieval, spending extra bit-width on the layers whose quantization moves the gap most recovers most of an extra bit's benefit for half its cost. In classification, routing the few low-gap inputs to FP recovers most of the lost accuracy at a fraction of the FP cost. We advocate that compression should be judged beyond standard metrics, and the open question is whether training can encourage the gap that a target task benefits from, making a model more robust to quantization under specific deployment needs.

Reproducibility statement

Every number in this paper comes from code we release, and every table and figure is generated from stored result files. The theory is self-contained. Lemma 1 and the statements it rests on are proved in Appendix C and checked by exhaustive enumeration (Appendix B, *Boundary cases and verification*). Section 4 gives the models, datasets, quantization settings and calibration protocol. All calibration is label-free, and the calibration protocol in Appendix A.2 states, per procedure, which calibration splits are disjoint from evaluation and where retrieval overlaps. The supplementary material contains the code and the stored result files needed to regenerate every table and figure in the paper.

Ethics statement

This work studies the reliability of an existing deployment practice and introduces no new models, data collection, or human subjects. All datasets are public, standard benchmarks used under their original licenses. The practical consequence we report is a safety one. Quantized retrieval systems can degrade in ways their accuracy metrics do not reveal, which may disadvantage users whose queries fall in the affected minority. The stability check and the allocation criterion are offered to make that risk measurable and reducible before deployment. We report the settings in which our stability check fails alongside those in which it holds.

Use of large language models

Large language models were used as a general-purpose coding and writing assistant: drafting and refactoring experiment and plotting scripts, and editing prose for concision. All research questions, experimental designs, analyses, and claims are the authors' own, and the authors verified every reported number against the stored result files.

Limitations

Our results cover weight quantization at 4 and 3 bits (W4, W3), and eight-bit activations on top of W4 and W8 with one static range per tensor. Finer activation schemes such as per-token ranges are untested. The downstream evidence of RAG damage is one pipeline: a Qwen3-Emb-0.6B retriever with a Qwen3-8B generator on Natural Questions at W4. The stability check assumes that calibration and deployment inputs come from the same distribution, the standard conformal assumption, and will degrade under distribution shift. Retrieval results use exact cosine similarity over a finite corpus. Approximate nearest-neighbor indexes, which add ranking errors of their own, are out of scope.

References

MNIST handwritten digit database, Yann LeCun, Corinna Cortes and Chris Burges. URL https://yann.lecun.com/exdb/mnist/. 21

Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefler, and James Hensman. Quarot: Outlier-free 4-bit inference in rotated llms. *Advances in Neural Information Processing Systems*, 37:100213–100240, 2024. 2

Hicham Badri and Appu Shaji. Half-quadratic quantization of large machine learning models, November 2023. URL https://mobiusml.github.io/hqq_blog/. 2, 7

Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. Ms marco: A human generated machine reading comprehension dataset, 2018. URL https://arxiv.org/abs/1611.09268. 4, 7

Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapid-deployment. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019. 28

Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – Mining Discriminative Components with Random Forests. In David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (eds.), *Computer Vision – ECCV 2014*, pp. 446–461, Cham, 2014. Springer International Publishing. ISBN 978-3-319-10599-4. doi: 10.1007/978-3-319-10599-4_29. 21

Vera Boteva, Demian Gholipour, Artem Sokolov, and Stefan Riezler. A full-text learning to rank dataset for medical information retrieval. In *European Conference on Information Retrieval*, pp. 716–722. Springer, 2016. 4

Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13166–13175. IEEE, 2020. 3

Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. Efficient intent detection with dual sentence encoders. In Tsung-Hsien Wen, Asli Celikyilmaz, Zhou Yu, Alexandros Papangelis, Mihail Eric, Anuj Kumar, Iñigo Casanueva, and Rushin Shah (eds.), *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pp. 38–45, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.nlp4convai-1.5. URL https://aclanthology.org/2020.nlp4convai-1.5. 21

Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote Sensing Image Scene Classification: Benchmark and State of the Art. *Proceedings of the IEEE*, 105(10):1865–1883, October 2017. ISSN 1558-2256. doi: 10.1109/JPROC.2017.2675998. URL https://ieeexplore.ieee.org/document/7891544/?arnumber=7891544. Conference Name: Proceedings of the IEEE. 21

Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. PACT: Parameterized clipping activation for quantized neural networks, 2018. URL https://openreview.net/forum?id=By5ugjyCb. 29

Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing Textures in the Wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3606–3613, Columbus, OH, USA, June 2014. IEEE. ISBN 978-1-4799-5118-5. doi: 10.1109/CVPR.2014.461. URL https://ieeexplore.ieee.org/document/6909856. 21

cjadams, Daniel Borkan, inversion, Jeffrey Sorensen, Lucas Dixon, Lucy Vasserman, and nithum. Jigsaw unintended bias in toxicity classification, 2019. URL https://kaggle.com/competitions/jigsaw-unintended-bias-in-toxicity-classification. 21

Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. Deep Learning for Classical Japanese Literature, November 2018. URL http://arxiv.org/abs/1812.01718. arXiv:1812.01718 [cs, stat]. 21

Adam Coates, Andrew Ng, and Honglak Lee. An Analysis of Single-Layer Networks in Unsupervised Feature Learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, pp. 215–223. JMLR Workshop and Conference Proceedings, June 2011. URL https://proceedings.mlr.press/v15/coates11a.html. ISSN: 1938-7228. 21

Arman Cohan, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel S Weld. Specter: Document-level representation learning using citation-informed transformers. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pp. 2270–2282, 2020. 4

Gregory Cohen, Saeed Afshar, Jonathan Tapson, and André van Schaik. EMNIST: Extending MNIST to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pp. 2921–2926, May 2017. doi: 10.1109/IJCNN.2017.7966217. URL https://ieeexplore.ieee.org/document/7966217/?arnumber=7966217. ISSN: 2161-4407. 21

Zhen Dong, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Hawq: Hessian aware quantization of neural networks with mixed-precision. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 293–302, 2019. 3

Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Hawq-v2: Hessian aware trace-weighted quantization of neural networks. *Advances in neural information processing systems*, 33:18518–18529, 2020. 3, 28

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 4

Abhinav Dutta, Sanjeev Krishnan, Nipun Kwatra, and Ramachandran Ramjee. Accuracy is not all you need. *Advances in Neural Information Processing Systems*, 37:124347–124390, 2024. 3

Peter C Fishburn. Intransitive indifference with unequal indifference intervals. *Journal of Mathematical Psychology*, 7(1):144–149, 1970. 2, 3

Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur, and Prem Natarajan. Massive: A 1m-example multilingual natural language understanding dataset with 51 typologically-diverse languages, 2022. 21

Elias Frantar and Dan Alistarh. Optimal brain compression: A framework for accurate post-training quantization and pruning. *Advances in Neural Information Processing Systems*, 35:4475–4488, 2022. 29

Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. GPTQ: Accurate post-training quantization for generative pre-trained transformers. In *International Conference on Learning Representations (ICLR)*, 2023. 2, 7, 9, 29

Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In Maria Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1050–1059, New York, New York, USA, 20–22 Jun 2016. PMLR. URL https://proceedings.mlr.press/v48/gal16.html. 9

Yonatan Geifman and Ran El-Yaniv. Selective classification for deep neural networks. *Advances in neural information processing systems*, 30, 2017. 3

Evangelos Georganas, Dhiraj Kalamkar, Alexander Kozlov, and Alexander Heinecke. Ml-specqd: Multi-level speculative decoding with quantized drafts, 2025. URL https://arxiv.org/abs/2503.13565. 3

Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. In *Low-power computer vision*, pp. 291–326. Chapman and Hall/CRC, 2022. 1, 2

Ian J. Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, Dong-Hyun Lee, Yingbo Zhou, Chetan Ramaiah, Fangxiang Feng, Ruifan Li, Xiaojie Wang, Dimitris Athanasakis, John Shawe-Taylor, Maxim Milakov, John Park, Radu Ionescu, Marius Popescu, Cristian Grozea, James Bergstra, Jingjing Xie, Lukasz Romaszko, Bing Xu, Zhang Chuang, and Yoshua Bengio. Challenges in Representation Learning: A Report on Three Machine Learning Contests. In Minho Lee, Akira Hirose, Zeng-Guang Hou, and Rhee Man Kil (eds.), *Neural Information Processing*, pp. 117–124, Berlin, Heidelberg, 2013. Springer. ISBN 978-3-642-42051-1. doi: 10.1007/978-3-642-42051-1_16. 21

Parvel Gu. Detecting a route flip is easier than knowing whether to fix it: Causal route-mediated damage in quantized mixture-of-experts. *arXiv preprint arXiv:2608.11212*, 2026. 2

Ruiqi Guo, Philip Sun, Erik Lindgren, Quan Geng, David Simcha, Felix Chern, and Sanjiv Kumar. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning*, pp. 3887–3896. PMLR, 2020. 2, 3

Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, July 2019. ISSN 2151-1535. doi: 10.1109/JSTARS.2019.2918242. URL https://ieeexplore.ieee.org/document/8736785/?arnumber=8736785. Conference Name: IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing. 21

Dan Hendrycks and Kevin Gimpel. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *Proceedings of International Conference on Learning Representations*, 2017. 3

Sara Hooker, Aaron Courville, Gregory Clark, Yann Dauphin, and Andrea Frome. What do compressed deep neural networks forget? *arXiv preprint arXiv:1911.05248*, 2019. 3

Sara Hooker, Nyalleng Moorosi, Gregory Clark, Samy Bengio, and Emily Denton. Characterising bias in compressed models. *arXiv preprint arXiv:2010.03058*, 2020. 3

Qiang Hu, Yuejun Guo, Maxime Cordy, Xiaofei Xie, Wei Ma, Mike Papadakis, and Yves Le Traon. Characterizing and understanding the behavior of quantized models for reliable deployment. *arXiv preprint arXiv:2204.04220*, 2022. 3

Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Weinberger. Multi-scale dense networks for resource efficient image classification. In *International Conference on Learning Representations*, 2018. 3

Gabriel Ilharco, Mitchell Wortsman, Samir Yitzhak Gadre, Shuran Song, Hannaneh Hajishirzi, Simon Kornblith, Ali Farhadi, and Ludwig Schmidt. Patching open-vocabulary models by interpolating weights. *Advances in Neural Information Processing Systems*, 35:29262–29277, 2022. 21

Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing models with task arithmetic. In *The Eleventh International Conference on Learning Representations*, 2023. 21

Hervé Jégou, Romain Tavenard, Matthijs Douze, and Laurent Amsaleg. Searching in one billion vectors: re-rank with source coding. In *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 861–864. IEEE, 2011. 2

Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-class active learning for image classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2372–2379, 2009. doi: 10.1109/CVPR.2009.5206627. 3

Phillip Keung, Yichao Lu, György Szarvas, and Noah A. Smith. The multilingual amazon reviews corpus, 2020. 21

OM Kiselev. Boundary-aware quantization: Finite-scale decision geometry of neural classifiers. *arXiv preprint arXiv:2607.01478*, 2026. 2

Steven Kolawole, Don Dennis, Ameet Talwalkar, and Virginia Smith. Revisiting cascaded ensembles for efficient inference. In *Workshop on Efficient Systems for Foundation Models II @ ICML2024*, 2024. 3

Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3D Object Representations for Fine-Grained Categorization. In *2013 IEEE International Conference on Computer Vision Workshops*, pp. 554–561, Sydney, Australia, December 2013. IEEE. ISBN 978-1-4799-3022-7. doi: 10.1109/ICCVW.2013.77. URL http://ieeexplore.ieee.org/document/6755945/. 21

Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009. URL https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf. 21

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466, 2019. doi: 10.1162/tacl_a_00276. URL https://aclanthology.org/Q19-1026/. 4, 7

Yann Le and Xuan Yang. Tiny imagenet visual recognition challenge. pp. 3, 2015. 21

Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523): 1094–1111, 2018. 4

Haoran Li, Abhinav Arora, Shuohui Chen, Anchit Gupta, Sonal Gupta, and Yashar Mehdad. MTOP: A comprehensive multilingual task-oriented semantic parsing benchmark. In Paola Merlo, Jorg Tiedemann, and Reut Tsarfaty (eds.), *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 2950–2962, Online, April 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.257. URL https://aclanthology.org/2021.eacl-main.257. 21

Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*, 2023a. 4, 7

Zhikai Li, Junrui Xiao, Lianwei Yang, and Qingyi Gu. Repq-vit: Scale reparameterization for post-training quantization of vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 17227–17236, 2023b. 2

Darryl D. Lin, Sachin S. Talathi, and V. Sreekanth Annapureddy. Fixed point quantization of deep convolutional networks. In *International Conference on Machine Learning (ICML)*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2849–2858, 2016. 29

Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of machine learning and systems*, 6: 87–100, 2024. 2, 7, 29

Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. Spinquant: Llm quantization with learned rotations. In *The Thirteenth International Conference on Learning Representations*, 2024. 2

R Duncan Luce. Semiorders and a theory of utility discrimination. *Econometrica, Journal of the Econometric Society*, pp. 178–191, 1956. 2, 3

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In Dekang Lin, Yuji Matsumoto, and Rada Mihalcea (eds.), *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL https://aclanthology.org/P11-1015. 21

Wei Chen Maggie, Phil Culliton. Tweet sentiment extraction, 2020. URL https://kaggle.com/competitions/tweet-sentiment-extraction. 21

Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk, and Alexandra Balahur. Www'18 open challenge: Financial opinion mining and question answering. In *Companion Proceedings of the The Web Conference 2018*, WWW '18, pp. 1941–1942, Republic and Canton of Geneva, CHE, 2018. International World Wide Web Conferences Steering Committee. ISBN 9781450356404. doi: 10.1145/3184558.3192301. URL https://doi.org/10.1145/3184558.3192301. 4

Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. Mteb: Massive text embedding benchmark. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 2014–2037, 2023. 4, 21

Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International conference on machine learning*, pp. 7197–7206. PMLR, 2020. 2, 7, 29

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2011*, 2011. URL http://ufldl.stanford.edu/housenumbers/nips2011_housenumbers.pdf. 21

Maria-Elena Nilsback and Andrew Zisserman. Automated Flower Classification over a Large Number of Classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pp. 722–729, December 2008. doi: 10.1109/ICVGIP.2008.47. URL https://ieeexplore.ieee.org/document/4756141. 21

James O'Neill, Polina Rozenshtein, Ryuichi Kiryo, Motoko Kubota, and Danushka Bollegala. I wish I would have loved this one, but I didn't – a multilingual dataset for counterfactual detection in product review. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 7092–7108, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.568. URL https://aclanthology.org/2021.emnlp-main.568. 21

Nilesh Prasad Pandey, Markus Nagel, Mart van Baalen, Yin Huang, Chirag Patel, and Tijmen Blankevoort. A practical mixed precision algorithm for post-training quantization, 2023. URL https://arxiv.org/abs/2302.05397. 29

Hancheol Park, Geonho Lee, Tairen Piao, and Tae-Ho Kim. Value-and-structure alignment for routing-consistent quantization of mixture-of-experts models. *arXiv preprint arXiv:2606.05688*, 2026. 2

Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3498–3505, June 2012. doi: 10.1109/CVPR.2012.6248092. URL https://ieeexplore.ieee.org/document/6248092. ISSN: 1063-6919. 21

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PmLR, 2021. 1, 4, 6

Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pp. 3982–3992, 2019. 21

Elvis Saravia, Hsien-Chi Toby Liu, Yen-Hao Huang, Junlin Wu, and Yi-Shin Chen. CARER: Contextualized affect representations for emotion recognition. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3687–3697, Brussels, Belgium, October-November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1404. URL https://aclanthology.org/D18-1404. 21

Dana Scott and Patrick Suppes. Foundational aspects of theories of measurement1. *The journal of symbolic logic*, 23(2):113–128, 1958. 2, 3

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In David Yarowsky, Timothy Baldwin, Anna Korhonen, Karen Livescu, and Steven Bethard (eds.), *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1631–1642, Seattle, Washington, USA, October 2013. Association for Computational Linguistics. URL https://aclanthology.org/D13-1170. 21

Mohamed A Soliman and Ihab F Ilyas. Ranking with uncertain scores. In *2009 IEEE 25th international conference on data engineering*, pp. 317–328. IEEE, 2009. 2

Johannes Stallkamp, Marc Schlipsing, Jan Salmen, and Christian Igel. The German Traffic Sign Recognition Benchmark: A multi-class classification competition. In *The 2011 International Joint Conference on Neural Networks*, pp. 1453–1460, July 2011. doi: 10.1109/IJCNN.2011.6033395. URL https://ieeexplore.ieee.org/document/6033395/?arnumber=6033395. ISSN: 2161-4407. 21

Edward Szpilrajn. Sur l'extension de l'ordre partiel. *Fundamenta mathematicae*, 16(1):386–389, 1930. 2, 3

Wei Tao, Haocheng Lu, Xiaoyang Qu, Bin Zhang, Kai Lu, Jiguang Wan, and Jianzong Wang. Moqae: Mixed-precision quantization for long-context llm inference via mixture of quantization-aware experts. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 10810–10820, 2025. 3

Surat Teerapittayanon and Bradley McDanel. Branchynet: Fast inference via early exiting from deep neural networks. In *2016 23rd international conference on pattern recognition (ICPR)*, pp. 2464–2469. IEEE, 2016. 3

Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. Beir: A heterogenous benchmark for zero-shot evaluation of information retrieval models. *arXiv preprint arXiv:2104.08663*, 2021. 4

Bastiaan S. Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation Equivariant CNNs for Digital Pathology. In Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pp. 210–218, Cham, 2018. Springer International Publishing. ISBN 978-3-030-00934-2. doi: 10.1007/978-3-030-00934-2_24. 21

Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*. Springer, 2005. 4

David Wadden, Shanchuan Lin, Kyle Lo, Lucy Lu Wang, Madeleine van Zuylen, Arman Cohan, and Hannaneh Hajishirzi. Fact or fiction: Verifying scientific claims. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pp. 7534–7550, 2020. 4

Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*, 2022. 4, 7

Ross Wightman. Pytorch image models. https://github.com/rwightman/pytorch-image-models, 2019. 21

Hao Wu, Patrick Judd, Xiaojie Zhang, Mikhail Isaev, and Paulius Micikevicius. Integer quantization for deep learning inference: Principles and empirical evaluation, 2020. URL https://arxiv.org/abs/2004.09602. 29

Zhuguanyu Wu, Jiaxin Chen, Hanwen Zhong, Di Huang, and Yunhong Wang. Adalog: Post-training quantization for vision transformers with adaptive logarithm quantizer. In *European Conference on Computer Vision*, pp. 411–427. Springer, 2024. 2

Guangxuan Xiao, Ji Lin, Mickael Seznec, Hao Wu, Julien Demouth, and Song Han. Smoothquant: Accurate and efficient post-training quantization for large language models. In *International conference on machine learning*, pp. 38087–38099. PMLR, 2023. 2

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms, September 2017. URL http://arxiv.org/abs/1708.07747. arXiv:1708.07747 [cs, stat]. 21

Jianxiong Xiao, Krista A. Ehinger, James Hays, Antonio Torralba, and Aude Oliva. SUN Database: Exploring a Large Collection of Scene Categories. *International Journal of Computer Vision*, 119(1):3–22, August 2016. ISSN 1573-1405. doi: 10.1007/s11263-014-0748-y. URL https://doi.org/10.1007/s11263-014-0748-y. 21

Shitao Xiao, Zheng Liu, Peitian Zhang, Niklas Muennighoff, Defu Lian, and Jian-Yun Nie. C-pack: Packed resources for general chinese embeddings. In *Proceedings of the 47th international ACM SIGIR conference on research and development in information retrieval*, pp. 641–649, 2024. 4, 7

Zhewei Yao, Zhen Dong, Zhangcheng Zheng, Amir Gholami, Jiali Yu, Eric Tan, Leyuan Wang, Qijing Huang, Yida Wang, Michael Mahoney, et al. Hawq-v3: Dyadic neural network quantization. In *International conference on machine learning*, pp. 11875–11886. PMLR, 2021. 3

Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the association for computational linguistics*, 2:67–78, 2014. 4, 6

Zhihang Yuan, Chenhao Xue, Yiqi Chen, Qiang Wu, and Guangyu Sun. Ptq4vit: Post-training quantization for vision transformers with twin uniform quantization. In *European conference on computer vision*, pp. 191–207. Springer, 2022. 2

Jingtao Zhan, Jiaxin Mao, Yiqun Liu, Jiafeng Guo, Min Zhang, and Shaoping Ma. Jointly optimizing query encoder and product quantization to improve retrieval performance. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pp. 2487–2496, 2021. 3

Yanzhao Zhang, Mingxin Li, Dingkun Long, Xin Zhang, Huan Lin, Baosong Yang, Pengjun Xie, An Yang, Dayiheng Liu, Junyang Lin, et al. Qwen3 embedding: Advancing text embedding and reranking through foundation models. *arXiv preprint arXiv:2506.05176*, 2025. 1, 4

Juntao Zhao, Wenhao Lu, Sheng Wang, Lingpeng Kong, and Chuan Wu. Qspec: Speculative decoding with complementary quantization schemes. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pp. 4779–4795, 2025. 3

Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients, 2018. URL https://arxiv.org/abs/1606.06160. 29

A Additional results and experimental detail

This appendix collects the material referenced from the main text: the tables cited by number (§A.1), the datasets, the retrieval construction and the full calibration protocol (§A.2), ranking damage beyond the top-1 (§A.3), corpus size and model scale (§A.5), the stability check against confidence baselines and the predicted flip curve (§A.6), the CLIP arm in full (§A.7), the allocation criteria we compare against (§A.8), the robustness of the allocation to budgets and calibration seeds (§A.9), and activation quantization (§A.10).

A.1 Tables referenced from the main text

A note on Table 4. At W3, ViT-L/14 loses a third of its Recall@1, and ViT-B/32 collapses from 50.1% to 9.4%. The smaller encoder is more fragile at both bit-widths.

Table 4: A deployed cross-modal retrieval system. CLIP text-to-image on Flickr30k, 4000 images, 2000 caption queries, zero-shot. The relevant image for a caption is the one it was written for. Recall@1 is the metric the system is judged by. Parenthesized values are relative changes in Recall@1. *Fraction of the queries whose correct image was ranked first under FP and no longer is.

Encoder	Bits	Top-1 change	R@1 (FP)	R@1 (PTQ)	Correct image lost*
ViT-B/32	W4	47.0%	50.1%	41.9% (−16.4%)	28.4%
ViT-B/32	W3	91.9%	50.1%	9.4% (−81.2%)	87.9%
ViT-L/14	W4	24.6%	59.8%	59.1% (−1.3%)	10.9%
ViT-L/14	W3	58.1%	59.8%	40.1% (−33.0%)	43.6%

Table 5: Among allocation criteria that need only forward passes, gap sensitivity is first under every quantizer. Fraction of the W3→W4 reduction in top-1 change recovered at a 3.5-bit average, pooled over models and corpora. Every criterion here is computable inside a serving deployment: no backward pass, no autograd, no calibration loss to choose.

Allocation criterion	RTN	HQQ	GPTQ	GPTQ (act-order)	AWQ	AdaRound
Gap sensitivity (ours)	**73.4%**	**72.8%**	**60.2%**	**63.2%**	**60.7%**	**69.4%**
Activation norm	55.0%	55.9%	36.8%	36.1%	40.7%	55.2%
Reconstruction error	47.5%	43.8%	44.2%	43.9%	46.6%	39.8%
Relative error (SQNR)	47.0%	47.7%	50.7%	51.6%	55.2%	38.9%
Depth heuristic	47.3%	41.3%	42.7%	43.6%	44.1%	35.8%
Activation-weighted error	44.3%	48.6%	30.5%	29.1%	38.4%	51.0%
Activation salience	42.8%	46.0%	30.2%	31.2%	38.7%	50.8%

Table 6: The same comparison on the deployed cross-modal system. CLIP text-to-image retrieval on Flickr30k (ViT-B/32 and ViT-L/14, zero-shot): fraction of the W3→W4 reduction in top-1 change recovered at a 3.5-bit average, pooled over both encoders, one calibration seed per quantizer. Calibration captions are held out of evaluation. Gap sensitivity leads every forward-only criterion under every rounding method. It wins 6 of 6 runs against reconstruction error.

Allocation criterion	RTN	GPTQ	AWQ
Gap sensitivity (ours)	**60.9%**	**37.9%**	**44.8%**
Activation norm	38.4%	27.5%	28.9%
Depth heuristic	37.6%	34.5%	27.4%
Relative error (SQNR)	39.4%	17.1%	18.8%
Activation-weighted error	38.4%	24.2%	27.9%
Activation salience	36.5%	33.3%	27.3%
Reconstruction error	37.8%	22.4%	15.6%

Table 7: The asymmetry survives the strictest control. One encoder (Qwen3-Emb-0.6B), one dataset, one quantization: the *same* embeddings of the test split are used twice, as classification through a linear probe fitted on full-precision training embeddings and then frozen, and as retrieval, by cosine nearest neighbor within the split. Only the task differs, so the gap between the columns cannot be attributed to a different checkpoint, dataset, or objective. Separation is the median $g_2/2\varepsilon$ of the corresponding scores.

Dataset	Bits	Top-1 change rate		Median separation	
		Classification	Retrieval	Classification	Retrieval
Banking77	W4	4.8%	**32.0%**	1.16	0.057
Banking77	W3	24.6%	**63.5%**	0.34	0.024
Emotion	W4	14.8%	**48.6%**	0.85	0.059
Emotion	W3	44.2%	**87.6%**	0.35	0.023
MTOPIntent	W4	6.4%	**43.1%**	1.39	0.048
MTOPIntent	W3	28.4%	**74.9%**	0.43	0.021

Table 8: The separation ratio orders models across architectures, not only within one. Top-1 retrieval change and median separation over the four text corpora. Six models, two architecture families, a $24\times$ parameter range. Across all 12 model–bit-width configurations the separation ratio predicts the change rate at Spearman -0.979 ($p = 3e-08$). Restricted to W3, -0.943 over six models.

			W4-group_128		W3-group_128	
Model	Family	Params (B)	Change	Sep.	Change	Sep.
Qwen3-Emb-0.6B	Decoder	0.6	33.4%	0.060	86.0%	0.022
Qwen3-Emb-4B	Decoder	4	24.4%	0.120	75.2%	0.032
Qwen3-Emb-8B	Decoder	8	21.3%	0.143	46.9%	0.054
BGE-large	Encoder	0.34	22.6%	0.094	54.3%	0.037
E5-large-v2	Encoder	0.34	29.5%	0.071	70.9%	0.028
GTE-large	Encoder	0.34	20.0%	0.122	44.8%	0.055

Table 9: A quantized cross-encoder rewrites the reranked order, and the separation ratio predicts by how much. A full-precision bi-encoder retrieves a fixed top-100 shortlist per query. Only the cross-encoder reranker (`ms-marco-MiniLM-L-6-v2`) is quantized. 500 queries per corpus. Flip is the fraction of queries whose reranked top-1 changes. Across the six configurations the separation ratio orders the flip rate at Spearman -0.943 ($p = 0.005$).

Corpus	Bits	Flip	Median sep. ratio
SciFact	W4	4.8%	1.299
SciFact	W3	12.2%	0.454
NFCorpus	W4	10.0%	0.783
NFCorpus	W3	17.6%	0.308
FiQA	W4	9.8%	0.591
FiQA	W3	27.0%	0.230

Table 10: Stronger quantizers reduce ranking damage and do not remove it. Top-1 retrieval change rate (%) at group_128, mean $\pm$ standard error over tasks.

Bits	Quantizer	ViT-B/16	ViT-L/16	Qwen3-Emb-0.6B	GTE	BGE	E5	Qwen3-Emb-4B	Qwen3-Emb-8B
W4	RTN	45.5 $\pm$2.5	42.1 $\pm$1.5	33.4 $\pm$4.1	20.0 $\pm$3.0	22.6 $\pm$3.1	29.5 $\pm$4.1	24.4 $\pm$3.2	21.3 $\pm$3.8
	HQQ	38.9 $\pm$2.6	33.5 $\pm$1.4	30.9 $\pm$3.4	17.3 $\pm$3.1	19.1 $\pm$2.6	22.0 $\pm$3.1	22.2 $\pm$3.1	18.2 $\pm$2.9
	GPTQ	33.1 $\pm$1.3	29.9 $\pm$1.5	23.1 $\pm$3.0	15.6 $\pm$2.5	15.5 $\pm$2.6	19.3 $\pm$3.4	17.3 $\pm$2.4	15.7 $\pm$3.0
	GPTQ (act-order)	30.6 $\pm$1.6	29.6 $\pm$1.3	22.5 $\pm$3.3	14.9 $\pm$2.7	14.4 $\pm$2.3	16.7 $\pm$2.7	18.4 $\pm$3.0	14.5 $\pm$3.0
	AWQ	43.6 $\pm$2.3	41.4 $\pm$1.2	29.1 $\pm$3.5	16.8 $\pm$2.8	17.1 $\pm$2.4	18.5 $\pm$2.1	21.0 $\pm$3.2	17.0 $\pm$2.9
	AdaRound	33.8 $\pm$1.5	32.1 $\pm$1.6	28.7 $\pm$3.1	17.9 $\pm$2.7	18.9 $\pm$2.7	23.7 $\pm$3.4	20.1 $\pm$2.7	18.2 $\pm$3.0
W3	RTN	80.9 $\pm$2.0	92.4 $\pm$0.9	86.0 $\pm$4.1	44.8 $\pm$4.8	54.3 $\pm$6.5	70.9 $\pm$9.4	75.2 $\pm$6.3	46.9 $\pm$6.4
	HQQ	67.2 $\pm$2.5	61.7 $\pm$1.4	59.7 $\pm$5.8	35.2 $\pm$4.1	43.3 $\pm$5.5	46.6 $\pm$6.8	44.5 $\pm$6.1	34.1 $\pm$5.8
	GPTQ	61.5 $\pm$1.8	57.8 $\pm$2.5	54.7 $\pm$5.7	31.7 $\pm$4.7	36.5 $\pm$4.8	39.9 $\pm$5.8	48.4 $\pm$6.3	41.5 $\pm$6.6
	GPTQ (act-order)	59.0 $\pm$2.2	55.4 $\pm$2.5	53.4 $\pm$6.1	31.4 $\pm$4.6	33.6 $\pm$4.6	38.9 $\pm$4.9	48.4 $\pm$6.3	41.0 $\pm$6.3
	AWQ	78.3 $\pm$1.9	89.5 $\pm$1.4	69.1 $\pm$5.7	34.5 $\pm$3.8	38.5 $\pm$4.3	40.2 $\pm$5.5	47.6 $\pm$6.7	37.2 $\pm$5.8
	AdaRound	62.5 $\pm$2.2	59.6 $\pm$1.9	65.3 $\pm$7.4	36.7 $\pm$4.5	45.4 $\pm$6.4	52.6 $\pm$7.1	48.2 $\pm$5.1	38.7 $\pm$6.7

Table 11: Separation, not change rate alone, shows which configurations cross the stability threshold. Median separation ratio $g_2/2\varepsilon$, top-1 change rate and quantized accuracy at group_128, both ViTs, six quantizers.

		ViT-B/16					ViT-L/16				
		Classification			Retrieval		Classification			Retrieval	
Quantizer	Bits	Sep	Change	Acc	Sep	Change	Sep	Change	Acc	Sep	Change
RTN	W4	3.18	6.2%	81.0%	0.017	45.5%	5.63	2.9%	88.9%	0.019	42.1%
HQQ	W4	4.61	5.7%	81.5%	0.019	38.9%	8.30	2.1%	89.2%	0.027	33.5%
GPTQ	W4	10.91	2.9%	82.0%	0.029	33.1%	14.05	1.4%	89.3%	0.044	29.9%
GPTQ (act-order)	W4	13.11	2.5%	82.0%	0.033	30.6%	15.76	1.3%	89.3%	0.047	29.6%
AWQ	W4	4.17	6.2%	81.3%	0.017	43.6%	5.42	3.0%	88.9%	0.019	41.4%
AdaRound	W4	8.06	3.4%	81.8%	0.027	33.8%	9.35	1.7%	89.3%	0.031	32.1%
RTN	W3	0.64	25.4%	69.6%	0.006	80.9%	0.51	42.1%	55.5%	0.005	92.4%
HQQ	W3	1.19	12.1%	78.6%	0.009	67.2%	2.03	6.0%	87.5%	0.009	61.7%
GPTQ	W3	3.06	7.0%	81.0%	0.012	61.5%	3.01	3.8%	88.6%	0.017	57.8%
GPTQ (act-order)	W3	4.07	6.2%	81.2%	0.013	59.0%	3.39	3.4%	88.8%	0.018	55.4%
AWQ	W3	0.81	19.6%	72.9%	0.007	78.3%	0.58	33.8%	63.1%	0.006	89.5%
AdaRound	W3	2.23	8.4%	80.4%	0.012	62.5%	2.48	4.8%	88.2%	0.013	59.6%

Table 12: Absolute retrieval metrics under each allocation. Mean over the 24-configuration headline grid (six models × four corpora, RTN, average 3.5 bits for the middle rows). Flip is disagreement with the full-precision top-1. Recall@1 and gold retention are against relevance judgements. The capture ratios of Table 5 are computed from these absolute values.

Configuration	Flip ↓	R@1 ↑	Gold lost ↓
All W3	61.0%	35.1%	40.4%
Reconstruction error	43.1%	46.0%	19.9%
Gap sensitivity	**33.3%**	**48.5%**	**13.8%**
All W4	23.3%	50.5%	8.5%

Gap sensitivity (ours)

Quantizer	Capture	Changed	R→W
RTN	73.3%	770	201
HQQ	70.5%	654	128
GPTQ	63.6%	542	86
GPTQ (AO)	77.6%	532	91
AWQ	66.6%	649	124
AdaRound	70.4%	629	113

Reconstruction error

Quantizer	Capture	Changed	R→W
RTN	59.0%	895	244
HQQ	57.0%	690	153
GPTQ	41.1%	600	114
GPTQ (AO)	53.1%	595	112
AWQ	51.2%	725	158
AdaRound	54.8%	641	122

Table 13: Gap allocation keeps the pipeline closer to FP under every quantizer. Natural Questions at a 3.5-bit average, against reconstruction error, the objective existing allocators optimize. The other criteria are compared in Figure 5. *Capture* is the share of the W3→W4 reduction in top-1 change recovered, over 3,324 queries on a 100,000-passage subsample. *Changed* counts, of the 2,255 questions that have a gold short answer, how many the pipeline answers differently from FP, and *R→W* how many of those turn from correct to incorrect, with a FP Qwen3-8B generator reading five passages. Gap sensitivity is better on every measure under every quantizer.

Table 14: Per model: gap-driven allocation wins on every one, and reconstruction error does not track gap damage. Capture of the W3→W4 flip-rate benefit at a 3.5-bit average under RTN, per model, pooled over the four corpora, and the Spearman correlation between per-layer reconstruction error and per-layer gap sensitivity, mean over all RTN allocation runs with the range across runs and the number of layers n.

Model	Reconstruction error	Gap sensitivity	Layers n	Mean ρ	Range across runs
Qwen3-Emb-0.6B	55.6%	**68.3%**	196	-0.037	$[-0.15, +0.03]$
Qwen3-Emb-4B	40.9%	**73.5%**	252	$+0.118$	$[+0.06, +0.18]$
Qwen3-Emb-8B	47.1%	**73.3%**	252	$+0.150$	$[+0.11, +0.21]$
BGE-large-en-v1.5	37.0%	**72.7%**	145	-0.101	$[-0.16, -0.07]$
E5-large-v2	57.8%	**84.5%**	145	-0.278	$[-0.33, -0.22]$
GTE-large	40.7%	**72.0%**	145	-0.101	$[-0.14, -0.07]$

Table 15: Routing score and threshold sensitivity. Fraction of inputs routed and fraction of the quantization accuracy loss recovered, at three percentiles of the routing score on an unlabeled validation slice, W4-channel, averaged over the eligible tasks per backbone, those whose FP →PTQ accuracy gap is at least 0.5 points (count in parentheses).

		10th pct.		25th pct.		50th pct.	
Backbone	Score	Routed	Recovered	Routed	Recovered	Routed	Recovered
ViT-B/16 (17)	Gap (ours)	9.8%	**62.3%**	24.9%	**84.8%**	50.8%	**96.6%**
	MSP	10.1%	58.4%	25.0%	79.5%	49.6%	92.7%
	Entropy	9.8%	48.3%	24.8%	69.0%	49.7%	83.6%
ViT-L/16 (16)	Gap (ours)	10.2%	**71.9%**	24.5%	**92.5%**	50.5%	**99.3%**
	MSP	9.6%	68.6%	24.5%	89.6%	49.9%	98.2%
	Entropy	9.5%	60.1%	24.3%	84.2%	49.6%	98.0%
Qwen3-Emb-0.6B (8)	Gap (ours)	10.1%	**75.7%**	25.0%	**87.8%**	50.7%	**97.9%**
	MSP	10.1%	74.8%	25.4%	88.2%	50.0%	98.3%
	Entropy	10.3%	73.0%	25.4%	88.1%	50.2%	98.3%

A.2 DATASETS, RETRIEVAL CONSTRUCTION AND CALIBRATION

Models and settings. The ViTs use the `orig_in21k` initialization via `timm` (Wightman, 2019), and the fine-tuning protocol follows the task-arithmetic literature (Ilharco et al., 2023; 2022). The classification head is excluded from quantization throughout, and the five quantizers other than RTN are evaluated at the group granularity. Every ranking is by cosine similarity between L2-normalized embeddings. The ViTs use the pooled representation just before the classification head. The text embedders use their own pooling through `sentence-transformers` (Reimers & Gurevych, 2019), which for Qwen3-Embedding is last-token pooling with the instruction prefix these models expect on queries. CLIP uses its image and text projections. CLIP is the OpenAI ViT-B/32 and ViT-L/14 release, loaded through `open_clip`: an image ViT (12 blocks of width 768 on 32-pixel patches, or 24 blocks of width 1024 on 14-pixel patches) and a 12-block text Transformer (width 512 or 768), each followed by a linear projection into a shared space of 512 or 768 dimensions. Ranking is by cosine similarity, and there is no fine-tuning. Every linear layer of both encoders is quantized.

Quantizers. All six are weight-only and quantize every linear layer. RTN rounds to the nearest grid point with a symmetric absmax scale per channel or per group. GPTQ uses block size 128 and 1% dampening of the Hessian diagonal, and its activation-ordered variant quantizes columns in decreasing order of that diagonal. AWQ searches the scaling exponent on a grid of 20 values. HQQ fits scale and zero-point under an ℓ_p loss with $p = 0.7$ by half-quadratic splitting, 20 iterations from a min-max initialization. AdaRound keeps the RTN scale and learns each weight's rounding direction with the rectified sigmoid of the original paper, 400 Adam steps at learning rate 0.03, regularizer weight 10, annealed from $\beta = 20$ to 2 after a 20% warm-up. Its reconstruction loss is computed exactly from the calibration second moment $X^\top X$ rather than from minibatches. GPTQ, AWQ and AdaRound calibrate on 512 documents and quantize layers sequentially, each seeing the outputs of the already quantized layers before it. HQQ and RTN are data-free.

Datasets. The 21 image classification tasks are Cars (Krause et al., 2013), CIFAR-10 and CIFAR-100 (Krizhevsky & Hinton, 2009), DTD (Cimpoi et al., 2014), EMNIST (Cohen et al., 2017), EuroSAT (Helber et al., 2019), FashionMNIST (Xiao et al., 2017), FER2013 (Goodfellow et al., 2013), Flowers102 (Nilsback & Zisserman, 2008), Food101 (Bossard et al., 2014), GTSRB (Stallkamp et al., 2011), KMNIST (Clanuwat et al., 2018), MNIST (noa), OxfordIIITPet (Parkhi et al., 2012), PCAM (Veeling et al., 2018), RenderedSST2 (Socher et al., 2013), RESISC45 (Cheng et al., 2017), STL10 (Coates et al., 2011), SUN397 (Xiao et al., 2016), SVHN (Netzer et al., 2011) and TinyImageNet (Le & Yang, 2015). The 11 text classification tasks come from MTEB (Muennighoff et al., 2023): AmazonCounterfactual (O'Neill et al., 2021), AmazonReviews (Keung et al., 2020), Banking77 (Casanueva et al., 2020), Emotion (Saravia et al., 2018), IMDB (Maas et al., 2011), MassiveIntent, MassiveScenario (FitzGerald et al., 2022), MTOPDomain, MTOPIntent (Li et al., 2021), ToxicConversations (cjadams et al., 2019) and TweetSentimentExtraction (Maggie, 2020).

Image self-retrieval. The image setting is self-retrieval: the corpus is the task's test split, queries are drawn from it, and the self-match is removed before ranking. It is a controlled setup, not a deployed system, and we use it on purpose. Because the encoder is the *same* fine-tuned checkpoint that produces the classification results, it isolates the one thing that changes: the classification and retrieval numbers in Table 1 differ only in how the output is read, not in the model, the data, the quantizer, or the training. No standard retrieval benchmark offers that control. The text setting provides what this setup cannot (real corpora, real queries, and relevance judgements), and the two are always reported separately, never pooled.

Calibration protocol. Six procedures take calibration data (the allocation criteria, GPTQ, AWQ, AdaRound, routing, and the stability check) and none of them takes labels. The allocation criteria are estimated from at most 128 queries, capped at a quarter of the query pool, or in classification from 4 batches. GPTQ, AWQ and AdaRound collect activation statistics from 512 documents. Routing sets its threshold on an unlabeled validation slice. The stability check calibrates its threshold on a conformal split. Calibration and evaluation are disjoint throughout. In classification, the calibration inputs come from the *train* split while evaluation is on test. In retrieval, the calibration queries are a seeded random slice held out of the query pool, and every rate we report is computed on the complement. Routing and the stability check likewise fix their thresholds on held-out slices and report

Table 16: A finer quantization grid helps both task families and separates them further. Top-1 change rate at each scale granularity, W4 and W3 pooled, paired over the 2,004 configurations measured at all three.

	Classification			Retrieval		
Quantizer	Per channel	Group 128	Group 64	Per channel	Group 128	Group 64
RTN	42.1%	18.1%	9.6%	69.9%	50.0%	45.2%
HQQ	23.1%	6.6%	5.5%	57.1%	37.1%	32.9%
GPTQ	8.3%	4.5%	3.9%	48.0%	33.3%	30.1%
GPTQ (act-order)	8.0%	4.1%	3.6%	46.9%	32.1%	28.9%
AWQ	37.4%	14.4%	8.0%	63.2%	40.9%	36.1%
AdaRound	11.9%	5.1%	4.5%	56.8%	37.1%	33.6%

Table 17: Quantizing the retriever changes answers that the benchmark does not report. Retrieval-augmented generation on Natural Questions, 2,255 questions over 2.68M passages, with the index embedded at FP, only the query encoder Qwen3-Emb-0.6B quantized, and a FP Qwen3-8B generator decoding greedily over five passages, one row per weight quantizer at W4

. The last three columns count questions.

Condition	Exact match	Token F1	Gold in context	Answers changed	Right→wrong	Wrong→right
No passages	15.1%	26.8%	–	–	–	–
FP retrieval	42.7%	58.1%	65.0%	–	–	–
PTQ retrieval, RTN	41.2%	56.4%	61.8%	731	106	70
PTQ retrieval, HQQ	42.0%	57.2%	62.4%	677	101	83
PTQ retrieval, GPTQ	41.7%	57.1%	63.9%	548	77	54
PTQ retrieval, GPTQ (AO)	42.7%	58.0%	64.5%	589	65	64
PTQ retrieval, AWQ	41.9%	57.4%	62.9%	675	90	70
PTQ retrieval, AdaRound	41.8%	57.0%	63.3%	631	82	60

on the remainder. One overlap remains, on the document side: GPTQ, AWQ and AdaRound draw their calibration documents from the indexed corpus, as is standard for post-training quantization. Nothing that is evaluated is selected using them, since the quantizer never sees a query.

One pair of embeddings serves the whole corpus-size sweep, since restricting the corpus only changes which scores are compared. The same pair also gives the query-only condition, in which the corpus is embedded offline in FP and only the query encoder is quantized. For each input, we record the top-K scores, with $K = 10$ for classification and 200 for retrieval. Retrieval needs the longer list because a corpus places many candidates within 2ε of the winner.

A.3 Ranking damage beyond the top-1

The protection is specific to the top-1. On classification, where accuracy looks nearly unchanged, the damage below the first position is already visible. The top-k part of Lemma 1 guarantees the top-k set when the gap between rank-k and rank-$k+1$ items exceeds 2ε. Table 18 gives that gap divided by 2ε at every rank. Between first and second place it is 2.4 to 5.4, so the top-1 cannot change. At every rank below, it falls to 0.5 or less. Cross-entropy separates the correct class from the rest and separates nothing else, so there is exactly one boundary the stability check can use, which is why the one we deploy (§3) covers the top-1 only. On models whose top-1 accuracy is preserved, the exact top-5 set survives for only 28 to 47% of inputs. Preserved accuracy is only about the top two items and tells nothing about the rest of the ranking.

Accuracy against ranking, in detail. Table 19 correlates, across tasks, the accuracy a configuration loses against the fraction of inputs whose top-5 set it changes. Within each quantization setting, the correlation is $+0.327$, about a tenth of the variance, and significant in only three of eight settings. In the deployed W4-group_128 setting, it is undetectable on either vision backbone ($+0.089$, $p = 0.72$ and $+0.230$, $p = 0.34$). Accuracy preservation, as reported, is not evidence of ranking preservation.

Retrieval quality against relevance judgements. Everywhere above, we measure agreement between the FP and quantized rankings, on purpose: it needs no relevance labels, and it answers the deployment question directly. It leaves open whether the changed results are *worse*. The text corpora come with relevance judgements, so we can answer that.

Table 18: The protection is specific to the argmax. Median separation at the k-cut relative to the perturbation, $(g_{k+1} - g_k)/2\varepsilon$, and the fraction of inputs whose exact top-k *set* is unchanged. W4-group_128.

	Separation at the k-cut				Exact top-k set kept			
Backbone	k=1	k=2	k=3	k=5	k=1	k=2	k=3	k=5
ViT-B/16	4.04	0.38	0.20	0.11	94%	76%	62%	45%
ViT-L/16	5.36	0.46	0.23	0.12	97%	81%	66%	47%
Qwen3-Emb	2.42	0.19	0.11	0.08	97%	62%	42%	28%

Table 19: Accuracy damage does not predict ranking damage. Spearman correlation across tasks between the accuracy a configuration loses and the fraction of inputs whose exact top-5 set it changes, computed *within* each configuration so bit-width and granularity are held fixed. Pooled across configurations the correlation is $+0.834$, but that compares settings poor at both against settings good at both. Removing the configuration effect leaves $+0.327$, and within the deployed W4-group_128 setting nothing is detectable. *significant at 0.05.

Model	Configuration	n	Spearman	p
ViT-B/16	W4-channel	19	$+0.174$	0.477
ViT-B/16	W4-group_128	19	$+0.089$	0.716
ViT-B/16	W3-channel	19	$+0.247$	0.307
ViT-B/16	W3-group_128	19	$+0.214$	0.379
ViT-L/16	W4-channel	19	$+0.461^*$	0.047
ViT-L/16	W4-group_128	19	$+0.230$	0.344
ViT-L/16	W3-channel	19	$+0.655^*$	0.002
ViT-L/16	W3-group_128	19	$+0.496^*$	0.031

Its change rates are computed over the judged queries only, those carrying relevance labels, and so differ slightly from the all-query rates reported elsewhere. At W4 a fifth to a third of top-1 results change for a 3.1–7.5% relative loss in nDCG@10, so the changes cost quality. The mean also hides the per-query outcome: of the flips whose FP top-1 was a judged-relevant document, 67% replace it with one that is not (Table 20). We repeat the distinction of §5 here because these numbers could be read the other way.

At W3, nDCG@10 falls by 19–69%, from a working retrieval system to one that returns mostly unrelated documents. At this bit-width the ranking is no longer usable. That is the practical meaning of two earlier findings: the stability check accepts nothing at W3, and no quantizer we tested repairs it.

A mean is the wrong statistic for this failure, and reporting one repeats the mistake this paper is about. The per-query question is instead: for how many queries does a document that was both relevant and inside the FP top-k fall out of the quantized top-k?

At the deployed setting, 8B at W4, where mean nDCG@10 falls by only 3.1%, **15.6% of the queries whose top result was relevant lose it**. One in six of those queries is harmed while the average moves by only three percent. At 0.6B the figures are 7.5% and 22.2%, and the lost fraction is within a point of these at $k = 5$ and $k = 10$. Averaging over queries understates per-query ranking damage, just as averaging over inputs understates it in accuracy.

Scale helps, and much more at W3 than at W4. Across a 13× range of model size the W4 loss falls from 7.5% to 3.1%, with diminishing returns. The W3 loss falls from 68.9% to 18.9%: at 0.6B three-bit retrieval is unusable, at 8B it is degraded but usable. Quality improves faster than the change rate falls (85.9% to 48.6% at W3), so larger models change fewer results, and a larger share of the changes they make are between documents of similar relevance.

A.4 Does a flip replace a relevant document?

Table 20 joins the top-1 flips of Table 1 against the BEIR relevance judgements. The corpus permutation and the query subsample are seeded, so the stored document indices identify the retrieved documents and the join needs no re-encoding. A query enters the table when it carries at least one judgement.

Table 20: A flip usually costs the relevant document, and the separation ratio predicts which queries are affected. Top-1 flips joined against the BEIR relevance judgements, over six text embedders, four corpora and six quantizers, both encoders quantized on the full corpus (73 runs at W4, 72 at W3). *Harm* is the share of all queries in the bucket whose relevant top-1 document is replaced by one that is not judged relevant. The text below reads the table.

	W4-group_128			W3-group_128		
Separation $g_2/2\varepsilon$	Queries	Flip	Harm	Queries	Flip	Harm
$[0,\ 0.05)$	16,508	53.0%	6.4%	28,690	69.7%	11.2%
$[0.05,\ 0.1)$	11,069	30.5%	3.9%	12,746	47.5%	11.4%
$[0.1,\ 0.2)$	14,175	13.9%	2.5%	10,657	25.6%	8.2%
$[0.2,\ 0.5)$	16,011	1.9%	0.4%	6,350	6.1%	3.4%
$[0.5,\ 1)$	5,715	0.0%	0.0%	954	0.2%	0.2%
≥ 1	2,097	0.0%	0.0%	–	–	–

Table 21: Corpus size matters only when it tightens the separation. Top-1 change rate at the smallest and largest corpus, W4-group_128, paired over queries (p from McNemar's test). The final columns show why: adding images from the same label set drives the separation ratio $g_2/2\varepsilon$ down, while adding documents to an already diverse corpus does not.

Backbone	Dataset	Corpus size		Change rate		Separation	
		Small	Large	Small	Large	Small	Large
Qwen3-Emb	FiQA	1,000	57,638	42.7%	41.9%	0.069	0.058
Qwen3-Emb	NFCorpus	1,000	3,633	36.1%	34.5%	0.067	0.062
Qwen3-Emb	SCIDOCS	1,000	25,657	31.8%	35.2%	0.087	0.051
Qwen3-Emb	SciFact	1,000	5,183	28.8%	22.0%*	0.091	0.132
ViT-B/16	CIFAR100	1,000	10,000	33.3%	46.0%*	0.048	0.013
ViT-B/16	Cars	1,000	8,041	43.6%	52.2%*	0.028	0.015
ViT-B/16	Food101	1,000	25,250	31.4%	49.7%*	0.060	0.017
ViT-B/16	SUN397	1,000	19,850	22.7%	39.4%*	0.145	0.041
ViT-B/16	TinyImageNet	1,000	10,000	25.1%	40.5%*	0.099	0.026

Pooled over the 6,578 W4 flips, 14.6% replace a relevant document by one that is not, 7.1% exchange one relevant document for another, 11.5% promote a relevant document that was not first before, and the remaining 66.8% move between documents that carry no judgement at all. The mean relevance of the served document therefore falls by 0.010 at W4 and by 0.092 at W3: losses and gains partly cancel, the same pattern as the CLIP result in §5, where 6.5% of queries turn from right to wrong while 5.7% turn the other way.

A.5 Corpus size and model scale

Corpus size, case by case. Every comparison in Table 21 is paired over the same queries. On the vision datasets, the separation ratio collapses as the corpus grows (`SUN397` $0.145 \rightarrow 0.041$), on `FiQA` it is flat, and on `SciFact`, the one dataset whose change rate falls, it rises ($0.091 \rightarrow 0.132$). On CLIP, ViT-B/32 is flat, and ViT-L/14 falls as the index grows, because the correct image enters it (Table 27).

The deployable configuration is not safe. A deployment stack embeds its corpus offline in FP and quantizes only the query encoder. Less of the system is quantized, so one might expect less damage. On the text corpora it helps modestly, reducing changes by 3–10 points, and on CLIP it helps most, from 47.0% to 27.1% on ViT-B/32 and from 24.6% to 15.3% on ViT-L/14 (Appendix A.7, *Corpus size*). On the vision backbones, it is often *worse*. On ViT-L/16 it is worse on all five datasets, by up to 6.5 points. Across every dataset, the top-1 result still changes for 18.9–53.6% of queries. Keeping the index exact does not stabilize the ranking, because the perturbation that reorders the candidates acts on the query that ranks them, which an exact index does nothing to cancel.

Table 22: The stability check against confidence scores under one protocol. ViT-B/16, W4 group_128, 21 tasks, mean over 20 calibration splits. Split conformal makes *any* score valid, so every row respects its bound. The generic construction thresholds the score at its $(1-\alpha)$ quantile among flipped calibration inputs, which bounds the joint event for any score. What separates the scores is how many inputs each accepts at a given level.

	$\alpha = 0.10$		$\alpha = 0.05$		$\alpha = 0.01$	
Score	Cover	Viol.	Cover	Viol.	Cover	Viol.
Our stability check	85.2%	1.47%	80.0%	0.72%	67.7%	0.10%
Margin, generic calibration	82.3%	0.94%	77.9%	0.54%	69.1%	0.15%
Max-softmax-prob, generic calibration	74.4%	1.29%	69.4%	0.81%	61.1%	0.33%

Table 23: The stored thresholds of the stability check. The value $\hat{\tau}$ that a test input's quantized top-1/top-2 gap must reach to be accepted, at each level α, W4 group_128, median over 20 random calibration splits. Logits for classification, cosine similarity for the retrieval rows. Coverage and violation rates are in Table 3.

Setting	$\alpha = 0.10$	$\alpha = 0.05$	$\alpha = 0.01$
Classification (ViT-B/16)	0.174	0.286	0.787
Text retrieval (Qwen3)	0.0854	0.0982	0.139
CLIP ViT-B/32	0.0571	0.065	0.0779
CLIP ViT-L/14	0.0454	0.0503	0.0629

Scale reduces the effect and does not remove it. One might expect the problem to affect only small models. Across a $13\times$ range of parameters in one model family (Tables 10 and 8), the top-1 retrieval change rate falls from 33.4% to 21.3% and the separation ratio roughly doubles. That is a real improvement, and it is in the direction the mechanism predicts: a larger model spreads its embeddings further apart relative to the same quantization noise. The largest model still changes one query in five, three to six times the classification rate of the 0.6B model (the only size in this family with fine-tuned classifiers), and its separation is still far below the threshold at which anything can pass. Scale, like a better quantizer, reduces the damage without removing it.

A.6 Calibrating the stability check, and the predicted flip curve

Table 3 reports two quantities at each level α: the *coverage*, the fraction of inputs the stability check accepts, and the *violation rate*, the fraction of accepted inputs whose top-1 changed nonetheless. Validity requires the second to fall below α, and usefulness requires the first to be large. The violation rate runs at least six times below α in every classification row, and no accepted retrieval query had its top-1 change at any level. Evaluating the top-1 part of Lemma 1 against an input's own ε would instead hold identically, and so test nothing.

Comparison with confidence scores. Split conformal calibration makes *any* score valid, so the fact that our stability check respects its bound says nothing about the theory. What separates scores is how many inputs each accepts at the same guarantee. Table 22 runs maximum softmax probability, the standard confidence score, through the same protocol, under the generic construction that makes any score valid (threshold at the $(1-\alpha)$ quantile of the score among flipped calibration inputs).

At $\alpha = 0.10$, our stability check accepts 85.2% against maximum softmax probability's 74.4%, 80.0 against 69.4 at $\alpha = 0.05$, and 67.7 against 61.1 at $\alpha = 0.01$, with violation rates far inside the level in every row. The margin itself, under the generic construction, sits between the two (82.3, 77.9 and 69.1%) and matches ours at the strictest level. So the gain over MSP comes from the quantity being measured, distance to the decision boundary rather than confidence. The stability check of §3 adds a little more by calibrating the perturbation of the score differences rather than a generic quantile.

The predicted curve. Proposition 2 turns the gap profile into a predicted flip probability. Table 24 fits one effective-scale parameter per group on the aggregate rate and reports the columns for each $|\mathcal{C}_\varepsilon|$ as predictions. The shape follows in both task families: for classification the fitted scale is 0.92 and 0.87, so the worst-case ε is already close to the perturbation scale that matters, and the prediction

Table 24: Proposition 2 predicts the shape as well as the level. Proposition 2 models the perturbation as uniform on $[-\varepsilon, \varepsilon]$. Real perturbations are not uniform, so we fit one interval half-width per group, the *effective perturbation scale* s, reported as a fraction of ε and chosen so the predicted *aggregate* flip rate matches the observed one. The columns for each $|\mathcal{C}_\varepsilon|$ are then predictions with no further tuning. For classification $s \approx 0.9$: the worst-case ε is already close to the operative scale. For retrieval s is 0.12–0.24, the inflation from taking a maximum over tens of thousands of candidates.

Group	s	$\lvert\mathcal{C}_\varepsilon\rvert$	Predicted	Measured
Classification, ViT-B/16	0.92	1	12.2%	15.6%
		2	19.9%	20.1%
		3	24.5%	23.5%
		4	30.0%	30.0%
Classification, ViT-L/16	0.87	1	11.5%	13.7%
		2	18.4%	17.0%
		3	23.1%	19.8%
		4	25.7%	23.7%
Retrieval, ViT-B/16	0.12	1–50	26.8%	28.9%
		51–120	38.9%	36.7%
		181–199	48.6%	49.1%
Retrieval, Qwen3-Emb	0.24	1–50	6.5%	4.1%
		181–199	36.7%	36.8%

Table 25: Contender set and separation on CLIP. Both encoders at group_128, RTN: top-1 change rate, median separation ratio $g_2/2\varepsilon$, fraction of queries whose contender set is empty, fraction of changed top-1 results taken by the full-precision runner-up, and fraction of queries whose top-5 set is unchanged. Counterpart of Figures 2 and 3.

Encoder	Bits	Change	Separation	$\mathcal{C}_\varepsilon = \emptyset$	From runner-up	Top-5 set kept
ViT-B/32	W4	47.0%	0.079	0.0%	33.3%	2.5%
ViT-B/32	W3	91.9%	0.037	0.0%	4.7%	0.0%
ViT-L/14	W4	24.6%	0.151	1.1%	49.3%	11.7%
ViT-L/14	W3	58.1%	0.062	0.0%	19.3%	0.5%

needs almost no tuning. For retrieval it is 0.12–0.24. This reflects the fact that ε is a maximum over tens of thousands of candidates, so the scale is not a free parameter.

Scope. The relation between $|\mathcal{C}_\varepsilon|$ and the change rate holds at W4, the regime in which the stability check is useful. It does *not* hold at W3, where retrieval remains 40–50 points more fragile than classification at every matched width. When the perturbation is as large as the score range itself, the number of contenders no longer limits the flips. The direction of the perturbation starts to matter, and the width does not capture that. This is not an artifact of the threshold. A valid bound must take the maximum over all candidates of that input, and the candidate set differs between classification and retrieval, so no single threshold can be both valid and comparable across the two. We therefore state the guarantee with the worst-case $\varepsilon = \|z^{\text{PTQ}} - z^{\text{FP}}\|_\infty$ and report that relation per bit-width.

A.7 THE CLIP RESULTS IN FULL

Every retrieval measurement reported on the text arm is repeated here on CLIP text-to-image retrieval (Flickr30k, ViT-B/32 and ViT-L/14, zero-shot), from the same stored runs as Tables 4 and 6. Each table names its text counterpart.

Mechanism. At W4 the contender set is empty for 1.1% of ViT-L/14 queries and for none of ViT-B/32, against 74.3% of classification inputs (§5.2). The FP runner-up takes the vacated first place in 49% and 33% of the changed results, the rest coming from deeper in the ranking, and the exact top-5 set survives for 12% and 3% of queries. The two encoders order by separation, as the six text embedders do: ViT-L/14 at a median ratio of 0.151 changes 24.6% of its top-1 results, ViT-B/32

Table 26: Separation and change rate on CLIP under six quantizers. Median separation ratio $g_2/2\varepsilon$, top-1 change rate and Recall@1 at group_128, both encoders quantized, all 2,000 queries. Counterpart of Table 11.

		ViT-B/32			ViT-L/14		
Quantizer	Bits	Sep	Change	R@1	Sep	Change	R@1
RTN	W4	0.079	47.0%	41.9%	0.151	24.6%	59.1%
HQQ	W4	0.101	34.4%	46.5%	0.176	21.2%	59.1%
GPTQ	W4	0.165	21.9%	48.9%	0.204	16.0%	60.5%
GPTQ (act-order)	W4	0.181	21.3%	48.6%	0.278	16.6%	60.0%
AWQ	W4	0.110	31.6%	47.8%	0.170	22.4%	60.2%
AdaRound	W4	0.125	28.1%	47.9%	0.197	17.4%	60.4%
RTN	W3	0.037	91.9%	9.4%	0.062	58.1%	40.1%
HQQ	W3	0.053	65.0%	32.4%	0.090	40.6%	52.9%
GPTQ	W3	0.069	52.4%	40.3%	0.114	33.1%	56.9%
GPTQ (act-order)	W3	0.076	47.8%	43.1%	0.116	33.5%	55.6%
AWQ	W3	0.047	74.1%	26.0%	0.074	42.7%	53.1%
AdaRound	W3	0.062	56.1%	37.5%	0.095	38.1%	55.0%

Table 27: Corpus size on CLIP. Top-1 change rate at the smallest and largest image index, RTN W4-group_128, paired over the same 2,000 queries (*: $p < 0.05$, McNemar's test), with the median separation ratio at each size. Nested seeded subsets of the index, as in Table 21.

Encoder	Dataset	Corpus size		Change rate		Separation	
		Small	Large	Small	Large	Small	Large
ViT-B/32	Flickr30k	500	4,000	47.4%	47.0%	0.084	0.079
ViT-L/14	Flickr30k	500	4,000	34.2%	24.6%*	0.127	0.151

at 0.079 changes 47.0%. The 30 CLIP configurations, under six quantizers and four corpus sizes, are the starred points of Figure 3, and they sit on the same curve as the other 1,098.

Stability check. Under the protocol of §3, the check accepts 15.7% of queries on ViT-L/14 at $\alpha = 0.10$ and 1.5% on ViT-B/32, with no violation in any split, alongside 11.3% on text retrieval and 85.2% on classification (Table 3).

What a changed result costs. Flickr30k pairs each caption with one image, so a changed top-1 either loses the correct image, gains it, or moves between two wrong ones. Pooled over all changed results at W4, 26% of ViT-L/14's and 30% of ViT-B/32's lose the correct image, and 23% and 13% gain it. On BEIR under the same pooling the figures are 14.6% and 11.5% (Table 20), but there two thirds of the changes move between unjudged documents, and no such unjudged mass exists on Flickr30k. At W3 the losing share rises to 45% and 48%. Losses and gains partly cancel in the mean, which is why Recall@1 moves by a point while a quarter of the results change (§5.3).

The five other quantizers. GPTQ, with or without activation ordering, removes a third to a half of the change rate on both encoders at both bit-widths, AdaRound three tenths to two fifths, and HQQ and AWQ a tenth to a third, the same ordering as on text (Tables 10 and 11). On ViT-L/14 at W4, Recall@1 under GPTQ is a point above FP while 16.0% of its top-1 results differ from it. All five raise the separation, GPTQ with activation ordering the most, and the largest value any of them reaches is 0.278, nearly four times below the stability threshold.

Corpus size. On ViT-B/32 the change rate is flat from 500 to 4,000 images (47.4% to 47.0%, $p = 0.84$). On ViT-L/14 it falls from 34.2% to 24.6% as the index grows and the separation rises from 0.127 to 0.151, the reverse of the ViT self-retrieval in Table 21. The reason is which image is on top. At the full index, the queries whose correct image is ranked first (59.8% of them on ViT-L/14) have a median separation of 0.252 against 0.065 for the rest, and change their top-1 for 10.9% against 45.1%. On ViT-B/32 the figures are 0.150 against 0.044 and 28.4% against 65.8%. At 500 images, the correct image is in the index for 13% of queries, so nearly every query ranks among wrong

Table 28: Gap-driven allocation on CLIP under each quantizer. Fraction of the W3→W4 benefit captured at a 3.5-bit average, pooled over both encoders, on the three metrics of Table 12.

Quantizer	Allocation	Flip	R@1	Gold kept
RTN	Reconstruction error	37.8%	47.3%	43.4%
	Gap sensitivity	**60.9%**	**74.3%**	**68.8%**
GPTQ	Reconstruction error	22.4%	23.4%	23.6%
	Gap sensitivity	**37.9%**	**32.8%**	**39.5%**
AWQ	Reconstruction error	15.6%	21.3%	18.6%
	Gap sensitivity	**44.8%**	**67.1%**	**55.0%**

Table 29: Absolute CLIP retrieval metrics under each allocation. Mean over both encoders, RTN, average 3.5 bits for the middle rows. Columns as in Table 12.

Configuration	Flip ↓	R@1 ↑	Gold lost ↓
All W3	74.8%	24.9%	65.6%
Reconstruction error	60.2%	37.0%	45.6%
Gap sensitivity	**51.2%**	**43.9%**	**33.9%**
All W4	36.1%	50.5%	19.6%

images, where the top-1 is a near-tie. Contrastive training separated the matching pair from the rest, not one wrong image from another. Quantizing only the query encoder, with the index embedded at FP, helps more here than on text, from 47.0% to 27.1% on ViT-B/32 and from 24.6% to 15.3% on ViT-L/14, and still leaves a sixth to a quarter of the results changing.

Allocation. Gap sensitivity captures 60.9% of the flip-rate benefit under RTN against 37.8% for reconstruction error, and leads on all three metrics under every quantizer (Table 28). In absolute terms, it changes 51.2% of top-1 results against 60.2%, reaches 43.9% Recall@1 against 37.0%, and loses 33.9% of correct images against 45.6% (Table 29). It wins on each encoder separately: 57.3% against 25.8% on ViT-B/32 and 65.6% against 53.9% on ViT-L/14 (Table 30). Reconstruction error correlates with gap sensitivity across layers at $+0.12$ on ViT-B/32 and $+0.34$ on ViT-L/14 (Table 30), higher than on the text embedders. The encoder with the higher correlation is also the one on which reconstruction error recovers more, 53.9% against 25.8%. On both, it stays behind gap sensitivity. The allocation is also stable across budgets and calibration seeds on CLIP (Tables 31 and 32).

A.8 Allocation criteria

Capture is pooled as $\sum \text{gain} / \sum \text{headroom}$ rather than averaged over per-configuration ratios, so that a configuration with almost nothing to gain does not dominate through a noisy ratio. The paragraphs below give the score and where it comes from. Every criterion scores each linear layer ℓ with a single number s_ℓ. Layers are then ranked by s_ℓ divided by their parameter count, since bit-width is paid per weight, and the wider bit-width goes to the highest-ranked layers until the budget is spent. Write W_ℓ for the layer's weight matrix, $\Delta_\ell = W_\ell - Q(W_\ell)$ for the change quantization makes to it, and x for the layer's input. Expectations over x are taken on the calibration set, in one forward pass that records, for each input channel i, the two moments $\mathbb{E}[x_i^2]$ and $\mathbb{E}|x_i|$.

Gap sensitivity (ours). Quantize layer ℓ alone, leave the rest at FP, and record how far the FP top-1/top-2 gap moves, as in §6.1. It is the only criterion that scores a layer by its effect on the quantity the ranking depends on rather than on an error norm. Needs calibration queries.

Reconstruction error. $s_\ell = \|\Delta_\ell\|^2 / |W_\ell|$, the mean squared change to the layer's weights. This is the objective weight-space allocators minimize: per-channel bit allocation at a fixed average bit-width solves exactly this problem (Banner et al., 2019), and it is the weight-space factor that HAWQ-v2 weights by the Hessian trace (Dong et al., 2020). It is the main text's baseline. Data-free.

Table 30: Per encoder on CLIP. Columns as in Table 14: capture of the W3→W4 flip-rate benefit at a 3.5-bit average under RTN, and the layer-level correlation between reconstruction error and gap sensitivity.

Model	Reconstruction error	Gap sensitivity	Layers n	Mean ρ	Range across runs
ViT-B/32	25.8%	**57.3%**	72	$+0.116$	$[+0.09, +0.12]$
ViT-L/14	53.9%	**65.6%**	108	$+0.342$	$[+0.33, +0.37]$

Table 31: The allocation holds across budgets. Capture of the W3→W4 benefit at three fractional bit budgets under RTN, calibration seed 2038, pooled over six text embedders and four corpora (Text) and over both CLIP encoders. At an integer average uniform is the only allocation, so the question is only posed between integers.

System	Budget	Allocation	Capture: flip	Capture: R@1	Capture: gold
Text	3.25 bits	Reconstruction error	28.4%	45.5%	42.5%
		Gap sensitivity	**46.4%**	**68.3%**	**62.1%**
	3.5 bits	Reconstruction error	47.5%	70.4%	64.3%
		Gap sensitivity	**73.4%**	**86.9%**	**83.2%**
	3.75 bits	Reconstruction error	63.6%	79.3%	76.5%
		Gap sensitivity	**89.8%**	**97.5%**	**94.4%**
CLIP	3.25 bits	Reconstruction error	15.7%	22.4%	18.9%
		Gap sensitivity	**32.4%**	**41.1%**	**37.9%**
	3.5 bits	Reconstruction error	37.8%	47.3%	43.4%
		Gap sensitivity	**60.9%**	**74.3%**	**68.8%**
	3.75 bits	Reconstruction error	69.0%	79.9%	74.7%
		Gap sensitivity	**83.5%**	**93.4%**	**87.9%**

Relative error (SQNR). $s_\ell = \|\Delta_\ell\|^2 / \|W_\ell\|^2$, the same quantity made scale-free, so a layer with large weights is not ranked highly merely for having large weights. Per-layer SQNR of this form is the classical basis for allocating fixed-point bit-widths across layers (Lin et al., 2016), and SQNR remains the sensitivity signal in recent post-training mixed-precision allocators (Pandey et al., 2023), which measure it at the network output rather than on the weights. Data-free.

Activation-weighted error. $s_\ell = \sum_i \mathbb{E}[x_i^2]\, \|\Delta_\ell[:, i]\|^2$, reconstruction error measured in activation space rather than weight space. It is the diagonal approximation of the layerwise proxy loss GPTQ (Frantar et al., 2023) minimizes, since that layer Hessian is $2X^\top X$ and its diagonal is proportional to $\mathbb{E}[x_i^2]$. The proxy predates GPTQ: Nagel et al. (2020) derive it from a second-order expansion of the task loss, and Frantar & Alistarh (2022) solve it in Optimal Brain Surgeon form. It needs calibration data and no backward pass.

Activation salience. $s_\ell = \sum_i \mathbb{E}|x_i|\, \|\Delta_\ell[:, i]\|^2$, the same form with the first absolute moment rather than the second. AWQ (Lin et al., 2024) protects channels with large activations by rescaling them. This criterion applies the premise per layer rather than per channel. Needs calibration data.

Activation norm. $s_\ell = \sum_i \mathbb{E}[x_i^2]$, the energy passing through the layer, ignoring what quantization does to its weights. We are not aware of prior work that ranks layers this way. It is included as a control to separate data-awareness from the rest, and tests whether knowing which layers see large activations is by itself enough to rank them. Needs calibration data.

Depth heuristic. $s_\ell = |d_\ell - \bar{d}| / \frac{1}{2}(d_{\max} - d_{\min})$, the layer's distance from the middle of the stack, so the first and last blocks rank highest. This is inspired by the common practice of keeping the first and last layers of a quantized network at higher precision, reported as more sensitive to quantization (Zhou et al., 2018; Choi et al., 2018; Wu et al., 2020). Data-free and weight-free.

Three of the six baselines therefore depend on the calibration set exactly as gap sensitivity does. In the calibration-size ablation of Table 33, every criterion is recomputed from the same calibration slice at each size, so a comparison at a small budget does not silently give the baselines a larger one.

Table 32: The allocation is stable across calibration draws (RTN). Fraction of layers receiving the same bit-width under different calibration seeds $\{101, 202, 2038\}$, over all corpora and budgets on the text arm and at a 3.5-bit average on CLIP. Chance agreement for this two-way split is approximately 50%.

System	Model	Seed pairs	Layer agreement
Text	Qwen3-Emb-0.6B	12	88.8%
	Qwen3-Emb-4B	12	90.0%
	Qwen3-Emb-8B	12	93.0%
	BGE-large-en-v1.5	12	93.0%
	E5-large-v2	12	90.7%
	GTE-large	12	92.9%
	All	72	**91.4%**
CLIP	ViT-B/32	3	90.7%
	ViT-L/14	3	87.7%
	All		**88.9%**

A.9 Allocation robustness: budgets and calibration seeds

Deployment. Every layer is entirely at 3 or at 4 bits, so the allocation needs no new low-level code: each layer runs with the ordinary 3-bit or 4-bit matrix-multiply routine, dispatched per layer.

Budget generality. The result does not depend on the 3.5-bit budget (Table 31): gap sensitivity is first at 3.25 and 3.75 bits per weight as well. Capture rises with the budget, as it must, since there are more bits to place, and the ordering is unchanged at every budget.

The allocation is stable. A sensitivity estimated from 128 calibration queries could be fitting that particular sample (Table 32). It is not: across 72 seed pairs, two calibration samples agree on the bit-width of 91.4% of layers, against roughly 50% by chance for a two-way split. The layers that decide a ranking are a property of the model, not of the sample used to find them.

Calibration size. Agreement with the 128-query allocation (Table 33) rises from 80.6% at 8 queries to 95.2% at 256, against 93.2% between two 128-query draws from different pools, so the allocation has converged by 64 queries. Capture rises with the sample, from 66.6% at 8 queries to 73.8% at 128, and is flat beyond. The three data-free criteria stay at 45–47% at every size, as they must, and the activation-based ones do not improve with more data. Gap sensitivity beats all six baselines in 13 of 24 runs at 8 queries and in 20 to 22 of 24 from 128. On CLIP the ordering is the same at every size on both encoders (Table 33).

A.10 Activation quantization

Every result in the main text quantizes the weights only. Here we also quantize the activations to eight bits at the input of every linear layer. Each activation tensor gets one fixed scale. We set the scale before deployment by running the model on the same calibration data the weight quantizers use, recording the largest value the tensor takes, and keeping that value fixed from then on.

CLIP is excluded because our method cannot reach all of its activations. We quantize an activation by attaching a small function to a layer that intercepts the layer's input just before the layer runs. This works whenever the layer is called in the normal way. In CLIP, the attention output projection is not called in the normal way. The attention implementation in PyTorch applies that layer's weights directly inside one combined function, so the layer itself is never called and our function is never reached. This affects 24 of the model's 72 linear layers. We could still quantize the weights of those layers, since weights are changed in place, but not their activations. A run would then quantize activations on two-thirds of the layers and leave the rest untouched, and reporting that as activation quantization would overstate what was done. Every other model in the paper calls all of its linear layers in the normal way, so all of them are covered.

Table 34 reports the 165 configurations measured at W4, W4A8 and W8A8, at group_128 under RTN. Three things hold. The asymmetry is unchanged: retrieval changes 6.6× more top-1 results

Table 33: Gap sensitivity leads at every calibration size. 3.5-bit average under RTN. Text rows pool six embedders on four BEIR corpora, 24 runs each, and CLIP rows are one run per cell. *Agreement* is the fraction of layers given the same bit-width as the 128-query gap allocation of the same system, with the query pool fixed so only the sample size varies. Chance for this two-way split is about 50%, and on text two 128-query allocations drawn from different pools agree on 93.2%. *Capture* is the share of the W3→W4 reduction in top-1 change recovered by each criterion, pooled as $\sum$gain/ $\sum$headroom, every criterion recomputed from the same calibration slice. The best per row is in bold. On text, gap sensitivity beats all six baselines in 13 of 24 runs at 8 queries and in 18 to 22 of 24 at larger sizes. SCIDOCS has 1000 queries and the calibration slice is capped at a quarter of the pool, so its largest arm uses 250 rather than 256.

			Capture by criterion						
System	Queries	Agreement	Gap	Act. norm	Recon.	SQNR	Depth	Act.-wt.	Salience
Text	8	80.6%	**66.6%**	51.4%	45.6%	45.8%	46.4%	42.4%	41.1%
	32	89.2%	**71.1%**	54.7%	45.7%	45.7%	46.5%	43.7%	40.9%
	64	92.9%	**72.4%**	52.6%	45.7%	45.7%	46.4%	43.1%	40.5%
	128	—	**73.8%**	52.5%	45.6%	45.6%	46.3%	42.8%	41.0%
	256	95.2%	**73.2%**	52.7%	45.7%	45.5%	46.6%	42.9%	40.7%
CLIP ViT-B/32	8	86.1%	**52.9%**	36.9%	26.2%	36.7%	35.8%	37.1%	29.7%
	32	91.7%	**53.4%**	36.9%	26.0%	36.3%	35.5%	37.2%	29.6%
	64	94.4%	**57.7%**	36.6%	25.8%	36.2%	35.2%	36.9%	29.5%
	128	—	**57.3%**	36.0%	25.8%	35.9%	35.2%	36.3%	29.0%
	256	94.4%	**54.4%**	35.3%	25.8%	35.8%	34.7%	35.4%	27.9%
CLIP ViT-L/14	8	77.8%	**64.0%**	42.9%	54.1%	44.1%	40.5%	41.7%	46.1%
	32	88.0%	**60.3%**	42.5%	54.1%	44.2%	40.8%	41.4%	45.9%
	64	89.8%	**65.3%**	41.9%	53.8%	43.6%	40.3%	41.2%	45.9%
	128	—	**65.6%**	41.6%	53.9%	44.1%	40.8%	41.3%	46.5%
	256	96.3%	**65.4%**	41.3%	53.6%	43.1%	41.0%	41.5%	46.4%

Table 34: Eight-bit activations on top of W4 and W8. Top-1 change rate, median separation ratio, and the fraction of configurations above the stability threshold, over the 165 configurations present in all three settings, group_128 under RTN. Every column is measured against FP.

Family		W4	W4A8	W8A8
Classification ($n = 53$)	Top-1 change	4.5%	8.7%	6.5%
	Median separation	3.63	1.80	3.04
	Above threshold	92%	75%	91%
Retrieval ($n = 112$)	Top-1 change	29.8%	60.1%	52.6%
	Median separation	0.093	0.031	0.035
	Above threshold	0%	0%	0%

than classification at W4, 6.9× at W4A8, and 8.1× at W8A8. Among the activation-quantized configurations alone, the separation ratio still predicts the change rate, at Spearman -0.90 over all 168 configurations at W4 and -0.93 over 165 at W8. The relation also carries over from one source of perturbation to the other: a line fitted on the weight-only configurations at four and three bits predicts the activation-quantized ones it never saw at $R^2 = 0.88$ in both cases, with the same slope to within 0.02. The perturbation added by eight-bit activations is comparable in size to the error from W4 and much larger than the error from W8, which is why W8A8 damages retrieval more than W4 alone.

The result does not depend on the weight quantizer. Table 35 repeats the comparison under each of the six weight quantizers, at W4 and at W8, with the activation scheme unchanged. Under every quantizer, adding eight-bit activations raises the classification change rate to between 7 and 9% at W4 and to 6 or 7% at W8, and the retrieval change rate to between 56 and 60% and to 52 or 53%. Hence, retrieval changes seven to eight times more top-1 results than classification in every case. The separation ratio predicts the change rate over the activation-quantized configurations alone at Spearman -0.90 or stronger under every quantizer and both bit-widths. The weight quantizer matters little once activations are quantized: at W4 alone the retrieval change rate spans 20 to 30% across

Table 35: Eight-bit activations under every weight quantizer. Top-1 change rate at W4 and W8, without and with eight-bit activations (W4A8, W8A8), for classification and retrieval, the retrieval-to-classification ratio with activations quantized, and Spearman between median separation ratio and change rate over the activation-quantized configurations alone. Over the same 165 configurations as Table 34, group_128.

	Four-bit weights						Eight-bit weights					
	Classification		Retrieval				Classification		Retrieval			
Quantizer	W4	W4A8	W4	W4A8	Ratio	Spearman	W8	W8A8	W8	W8A8	Ratio	Spearman
RTN	4.5%	8.7%	29.8%	60.1%	6.9×	−0.90	0.3%	6.5%	2.5%	52.6%	8.1×	−0.93
HQQ	3.9%	8.2%	24.9%	57.4%	7.0×	−0.91	0.2%	6.5%	2.3%	52.5%	8.0×	−0.93
GPTQ	2.4%	7.4%	20.9%	57.1%	7.7×	−0.92	0.1%	6.4%	1.9%	52.6%	8.2×	−0.93
GPTQ (AO)	2.2%	7.2%	19.8%	56.5%	7.8×	−0.93	0.1%	6.5%	1.8%	52.9%	8.1×	−0.93
AWQ	4.4%	8.3%	25.6%	57.6%	7.0×	−0.93	0.2%	6.5%	2.3%	52.5%	8.1×	−0.94
AdaRound	2.8%	7.6%	23.8%	58.1%	7.6×	−0.90	0.2%	6.5%	2.2%	52.4%	8.1×	−0.94

quantizers, but with activations added, it spans 56 to 60%, and at W8, the six quantizers agree to within half a point. The activation error is then the larger of the two, and the error-minimizing weight quantizers lose their advantage.

Our gap-sensitivity allocation also survives activation quantization. We re-evaluate the stored allocations of the 24 text retrieval configurations, six embedders on four BEIR corpora, with eight-bit activation quantization, under every policy, so the allocation still assigns weight bits and the activation error is a fixed extra perturbation. Table 36 gives the average capture. Gap sensitivity drops from 73% to 67% and stays first in 23 of 24 configurations, while reconstruction error drops from 48% to 37% and again does not outperform random bit-width allocation.

B THEORY: DEFERRED STATEMENTS AND REMARKS

Why indifference is not transitive. Two candidates are indifferent when neither is certainly above the other. Certain dominance itself is transitive, because gaps add: if $z_i - z_j \geq 2\varepsilon$ and $z_j - z_k \geq 2\varepsilon$ then $z_i - z_k \geq 4\varepsilon$. Indifference is not, because a chain of candidates each within 2ε of the next can span an arbitrarily large total gap. This is why a candidate sitting far below the top-1 can still reach the first position, which §5.2 measures.

Boundary cases and verification. The definition of certain dominance uses $\geq$ rather than $>$, and the choice matters at the boundary. When $z_i - z_j$ is exactly 2ε, no perturbation of size ε can put j strictly above i, because the most it can do is bring the two level. The pair is therefore genuinely ordered and belongs in the relation. The proof of Proposition 1 rests on the equivalent statement that a ranking is realizable iff $z_{c_k} - z_{c_j} < 2\varepsilon$ for all positions $j < k$. We also checked it by exhaustive enumeration over 1.3M permutations in exact rational arithmetic, covering both generic score vectors and ones with gaps sitting exactly at 2ε, with no mismatches.

Why each score is limited separately. Proposition 1 assumes that no single score moves by more than ε, which is what $\|\delta\|_\infty \leq \varepsilon$ says, and the characterization depends on the limit having that form. Suppose instead the limit applied to the perturbation as a whole, as $\|\delta\|_2 \leq \varepsilon$ does. Concentrating δ on a single score then reaches rankings that spreading it across many cannot, so whether a ranking is reachable depends on how many pairs it reverses at once, and no single threshold on a pair can describe the reachable set. Weight quantization rounds each weight on its own, and the error it puts on each score is bounded separately, so the per-score limit matches what quantization does.

Why the ratio and not the gap. The scores in the two task families are not measured in the same units. A logit is unbounded and a cosine similarity lies in $[-1, 1]$, so a reader may worry that comparing ε across the two settings compares nothing. We never do. Every comparison is on the separation ratio $g_2/2\varepsilon$, which is dimensionless and invariant under any positive affine change of the scores. Replacing z by $cz + b$ with $c > 0$ scales both g_2 and ε by c and cancels b in both, so the ratio is unchanged, and so is the ranking and therefore every flip we count. Temperature scaling of logits is the case $c = 1/\tau$. The invariance does not survive a nonlinear map: on $z = (3, 2, 1)$ with $\delta = (0, 0.4, 0)$, the ratio is 1.25 on the logits and 2.59 after a softmax, which is a further reason to measure the gap on the scores a system ranks by rather than on probabilities.

Table 36: Gap allocation still leads when activations are quantized too. Share of the W3→W4 reduction in top-1 change recovered at a 3.5-bit average, pooled over the 24 text retrieval configurations, six embedders on four BEIR corpora, with and without eight-bit activations, under RTN. The same stored allocations are evaluated in both columns.

Allocation criterion	Weights only	Weights and activations
Gap sensitivity (ours)	**73.4%**	**67.0%**
Reconstruction error	47.5%	37.0%
Random split (mean of 3)	49.3%	37.7%

The general characterization. The whole ranking survives every perturbation exactly when $\mathcal{R}_\varepsilon(z) = \{\pi(z)\}$, and Lemma 1 is the case of the first position and of the first k.

Proposition 1 (Characterization of reachable rankings)**.** *$\sigma = (c_1, \ldots, c_N) \in \mathcal{R}_\varepsilon(z)$ iff $z_{c_k} - z_{c_j} < 2\varepsilon$ for all positions $j < k$.*

Both parts of Lemma 1 follow by applying this to one pair. We use only the top-1 part, but the general statement is about a whole ranking rather than one comparison, which is what the non-transitivity above and §5.2 rest on.

From reachability to probability. Proposition 1 says which rankings are reachable, not how often one is realized, and on retrieval the resulting bound says almost nothing. Almost every input has a nonempty contender set, so the statement that the ranking may change is true and useless. The gap profile holds more information than the bound uses, and a short calculation extracts it.

Proposition 2 (Flip probability)**.** *Let $\delta_1, \ldots, \delta_N$ be i.i.d. uniform on $[-\varepsilon, \varepsilon]$ and write the separation ratio $a_j = g_j / 2\varepsilon$. Then the top-1 survives with probability*

$$\Pr[\textit{top-1 stable}] \;=\; \int_0^1 \prod_{j \in \mathcal{C}_\varepsilon} \min(u + a_j,\, 1)\; du.$$

Two extreme cases check the formula. Write $w = |\mathcal{C}_\varepsilon|$ for the number of contenders. If every contender is tied with the winner, so that $a_j = 0$ for all of them, the integral becomes $\int_0^1 u^w du = 1/(w+1)$. That is what you would predict by treating the winner and its w contenders as interchangeable and asking for the chance that the winner comes out on top, so the formula agrees with the naive count in the case where the naive count should apply. At the other extreme, if every contender sits exactly at the boundary, so that $a_j = 1$, the integral is 1 and the flip probability is zero, which is Lemma 1 again. Real inputs fall between the two, and the interchangeable estimate systematically predicts too many flips there, because contenders sit at different distances from the winner rather than all at zero. Proposition 2 accounts for those distances.

Table 24 evaluates the prediction. Proposition 2 assumes the perturbations are independent and uniform, which is a modeling choice rather than something quantization guarantees. Whether it captures the shape of the measured curve is an empirical question.

Which score to route on. The routing rule of §6.2 thresholds the quantized gap. The following says that, up to order-equivalence, no other score matches the flippable sets at every budget.

Proposition 3 (Identifiability of the gap)**.** *Let $\mathcal{S}$ be the class of measurable scores $s(x) = g(z_{\mathrm{q}}(x))$ of the quantized logit vector $z_{\mathrm{q}}(x)$, and assume the gap g_{q} has an atom-free distribution. Then s admits, for every $\varepsilon > 0$, a threshold θ_ε such that $\{x : s(x) \le \theta_\varepsilon\}$ equals the ε-flippable set up to a null set, if and only if s preserves the strict order of g_{q} on almost every pair of inputs. In particular g_{q} itself qualifies, with threshold 2ε, and any score that re-ranks a positive-measure set of pairs against it fails at some budget.*

This is an identifiability statement, not an optimality one: once Lemma 1 fixes the flippable sets, only order-equivalents of the gap can match them at every budget. It concerns the worst-case flippable set at each budget, not which inputs a given perturbation actually changes. The rule of §6.2 uses one threshold rather than every budget, so the support for using the gap there is empirical, not this proposition. Proof in Appendix C.

C PROOFS

Full proofs of the statements in Section 3 and Appendix B.

Proof of Proposition 1. Write P_ε for the relation $i \prec j$ iff $z_i - z_j \geq 2\varepsilon$, and call a permutation a linear extension of P_ε if it places i before j whenever $i \prec j$. *Necessity.* Suppose δ realizes π and $i \prec j$ in P_ε, i.e. $z_i - z_j \geq 2\varepsilon$. Note $i \neq j$, since $\varepsilon > 0$ makes P_ε irreflexive. Then

$$(z_i + \delta_i) - (z_j + \delta_j) \;\geq\; (z_i - z_j) - |\delta_i| - |\delta_j| \;\geq\; 0.$$

The perturbed values are strictly ordered along π, hence pairwise distinct, so the inequality is strict and i precedes j. Thus π is a linear extension.

Sufficiency. Let π be a linear extension. For $j < k$, if $z_{c_k} - z_{c_j} \geq 2\varepsilon$ then $c_k \prec c_j$ in P_ε and π would have to place c_k first, a contradiction, so $z_{c_k} - z_{c_j} < 2\varepsilon$ for all $j < k$. Set

$$t_1 = z_{c_1} + \varepsilon, \qquad t_k = \min(z_{c_k} + \varepsilon,\; t_{k-1} - \eta)\,,$$

which unrolls to $t_k = \min_{j \leq k} \big(z_{c_j} + \varepsilon - (k-j)\eta\big)$. Then t is strictly decreasing and $t_k \leq z_{c_k} + \varepsilon$ by construction. The remaining requirement $t_k \geq z_{c_k} - \varepsilon$ holds provided $z_{c_k} - z_{c_j} \leq 2\varepsilon - (k-j)\eta$ for all $j \leq k$. The instance $j = k$ reads $0 \leq 2\varepsilon$ and is trivial. Choosing

$$0 \;<\; \eta \;<\; \min_{j<k} \frac{2\varepsilon - (z_{c_k} - z_{c_j})}{k-j},$$

a minimum over a nonempty finite set of strictly positive reals, makes this hold. Then $\delta_{c_k} = t_k - z_{c_k}$ satisfies $\|\delta\|_\infty \leq \varepsilon$ and realizes π.

Proof of Lemma 1. *Top-1.* Apply Proposition 1 with the pair $(z_{(1)}, z_{(2)})$: if $z_{(1)} - z_{(2)} \geq 2\varepsilon$, then no linear extension places any other candidate first, so every realizable ranking has the same leader. Conversely, if the gap is below 2ε, the transposition swapping the top two is a linear extension, hence realizable by Proposition 1.

Top-k. If $z_{(k)} - z_{(k+1)} \geq 2\varepsilon$ then every candidate at rank $\leq k$ certainly dominates every candidate at rank $> k$, so by Proposition 1 no realizable ranking moves one across the cut and the top-k set is invariant. If the gap is below 2ε, ranks k and $k+1$ are incomparable, so the transposition exchanging them extends to a linear extension and changes the set.

Proof of Corollary 1. Candidate j overtakes a exactly when $z_a + \delta_a < z_j + \delta_j$, i.e. when $\delta_j - \delta_a > z_a - z_j$. The top-1 survives iff this fails for every j, an event that depends on δ only through the differences $\delta_j - \delta_a$. For the second claim suppose the top-1 changed, so $a' \neq a$ with $z_a > z_{a'}$ and $z_{a'} + \delta_{a'} \geq z_a + \delta_a$. Let b' be the runner-up of $z + \delta$. Since $z_a + \delta_a \leq z_{b'} + \delta_{b'}$,

$$g_2(z+\delta) = (z_{a'} + \delta_{a'}) - (z_{b'} + \delta_{b'}) \leq (z_{a'} - z_a) + (\delta_{a'} - \delta_a) < \delta_{a'} - \delta_a \leq S,$$

using $z_{a'} - z_a < 0$. Finally, $\delta_{a'} - \delta_j \leq 2\|\delta\|_\infty$ for every j, so $S \leq 2\|\delta\|_\infty$, with equality only when two coordinates are perturbed by ε in opposite directions. Calibrating S directly avoids that loss. The single pair of top-1 and runner-up does *not* suffice: a third candidate can overtake a while $z_a - z_b$ is unchanged, which is why S maximizes over all challengers.

Proof of Proposition 3. *Necessity.* Suppose $g_{\mathrm{q}}(x) < g_{\mathrm{q}}(x')$. Pick a rational $c \in (g_{\mathrm{q}}(x), g_{\mathrm{q}}(x'))$ and apply the hypothesis at $\varepsilon = c/2$: there is a threshold θ with $\{z : s(z) \leq \theta\} = \{z : g_{\mathrm{q}}(z) < c\}$ up to a null set N_c. Outside N_c, x lies in the right-hand set so $s(x) \leq \theta$, and x' does not, so $s(x') > \theta$. Hence $s(x) < s(x')$. The union of exception sets over rational c is null.

Sufficiency. Fix $\varepsilon > 0$ and $c = 2\varepsilon$, and let $A = \{g_{\mathrm{q}} < c\}$, $B = \{g_{\mathrm{q}} > c\}$, and $\theta = \operatorname{ess\,sup}_A s$. Then $A \subseteq \{s \leq \theta\}$ up to a null set. If $B_0 = \{x' \in B : s(x') \leq \theta\}$ had positive measure, order preservation would force $s = \theta$ almost everywhere on B_0, hence (applying order preservation within $B_0 \times B_0$) g_{q} constant on B_0 up to a null set, contradicting atom-freeness. Since $\{g_{\mathrm{q}} = c\}$ is null by atom-freeness, $\{s \leq \theta\}$ equals the ε-flippable set up to a null set. The flippable set itself equals $\{g_{\mathrm{q}} < 2\varepsilon\}$ by Lemma 1 and its matching converse (the two-coordinate perturbation $z_i \mapsto z_i - \varepsilon$, $z_j \mapsto z_j + \varepsilon$ flips any input with $g_{\mathrm{q}} < 2\varepsilon$). □

Remark. Order preservation, not a pointwise factorization $s = h \circ g_{\mathrm{q}}$, is the exact condition threshold routing needs: it makes the bottom-p set of s coincide with the bottom-p set of g_{q} for every budget p, up to null sets.

Proof of Proposition 2. The top-1 is displaced by j iff $\delta_j - \delta_1 \geq g_j$, so it survives iff $\delta_j < \delta_1 + g_j$ for every j. Conditioning on δ_1 and using independence, $\Pr[\text{stable} \mid \delta_1] = \prod_j \Pr[\delta_j < \delta_1 + g_j]$. For δ_j uniform on $[-\varepsilon, \varepsilon]$, $\Pr[\delta_j < t] = \min\big((t+\varepsilon)/2\varepsilon,\, 1\big)$ truncated below at 0. Substituting $u = (\delta_1 + \varepsilon)/2\varepsilon$, which is uniform on $[0,1]$, gives $\Pr[\delta_j < \delta_1 + g_j] = \min(u + a_j, 1)$, and integrating over u yields the claim. Candidates outside $\mathcal{C}_\varepsilon$ have $a_j \geq 1$ and contribute a factor of one.