

TWEO: Transformers Without Extreme Outliers Enables FP8 Training And Quantization For Dummies

Guang Liang^{1,2,3} Jie Shao^{1,2} Ningyuan Tang^{1,2} Xinyao Liu⁴ Jianxin Wu^{1,2*}

¹State Key Laboratory for Novel Software Technology, Nanjing University, China

²School of Artificial Intelligence, Nanjing University, China

³Zhongguancun Academy, Beijing, China

⁴University of Science and Technology of China, Hefei, China

{liangg, shaoj, tangny}@lamda.nju.edu.cn, liuxinyao@mail.ustc.edu.cn, wujx2001@nju.edu.cn

Abstract

*Native FP8 support in modern hardware is essential for training large Transformers, but is severely hindered by extreme activation outliers. Existing solutions either rely on complex mixed-precision engineering or invasive architectural modifications. This paper fundamentally challenges the conventional wisdom that outliers are data-driven. We demonstrate that extreme outliers are a data-independent, mechanically-produced artifact of training, originating from specific structural properties of the weight matrices (i.e., collinearity). Based on this insight, we propose TWEO (Transformers Without Extreme Outliers), a novel, non-invasive loss function. TWEO effectively prevents extreme outliers via a very simple loss term, which reduces outliers from 10000+ to less than 20. TWEO then enables **full-model FP8 pre-training** with neither engineering tricks nor architectural changes for both LLM and ViT. When standard FP8 training catastrophically collapses, TWEO achieves performance comparable to the BF16 baseline while delivering a 36% increase in training throughput. Also, TWEO enables a new quantization paradigm. Hardware-friendly **W&A8 per-tensor static quantization** of LLMs, previously considered completely unusable due to outliers, achieves SOTA performance for the first time on TWEO-trained models.*

1. Introduction

Computational demands of AI are growing at an unprecedented rate. In response, modern chips have introduced native support for 8-bit floating-point (FP8) computation, supported by software libraries such as Transformer Engine [22], which aims to double training throughput while significantly reducing memory bandwidth requirements.

But, it is still restricted by a core algorithmic barrier: the emergence of activation outliers, which are activation values whose magnitudes far exceed the average level [1, 10].

Outliers are often categorized as “normal” and “massive” [18, 43, 44]. Normal outliers (e.g., around 10) may be unavoidable from multiplication of two high-dimensional vectors, which does not pose challenges for FP8 training or quantization. We focus on massive or extreme outliers (e.g., >1000) [12], which pose serious challenges.

First, in Post-Training Quantization (PTQ), outliers are the primary obstacle [4, 53]. Extreme outliers drastically stretch the quantization range, forcing algorithms into a poor trade-off between clipping and rounding errors, which leads to severe accuracy drop [34]. Experiments showed that simply zeroing out a few (0.1%) outliers causes a catastrophic 600-1000% spike in validation perplexity [12].

Second, in native low-bit training, they lead to catastrophic stability collapse. The range of FP8 is extremely limited (e.g., E4M3 ranges within ± 448) [22, 32]. When extreme outliers (>1000) appear during training, numerical overflows are unavoidable and the failure mode (training loss suddenly explodes) will appear [28]. This instability forces tricks such as reverting to BF16 or even FP32 for many sensitive components [36], which complicates the model design and largely removes FP8’s efficiency gain.

On why extreme outliers emerge, existing research mostly linked them to properties of the input data, such as token frequency [35], special tokens [8], or sticky tokens [7]. Another data-dependent perspective is the no-op hypothesis [4], which posits that the attention module attempts to learn a “no-op”, which pushes the corresponding logits towards negative infinity, and hence requires the preceding layer to output very high magnitudes. More recent analyses connect massive activations and sink phenomena to architectural and normalization-driven rescaling effects in attention and residual streams [38, 44].

*Corresponding author.

In this paper, we first identify the cause of extreme outliers, then design a method to eliminate them. Following [42], we design a *Contradiction Stethoscope* to test whether outliers are fundamentally tied to the semantics or statistics of the input data. Our evidence shows that extreme outliers can be reproduced even under random inputs, and disappear when the trained weights are removed. This indicates that the *root cause* is mostly data-independent, stemming from the interaction of network architecture, optimization algorithms, and training dynamics. In vision, image patches do not have an equivalent to the high-frequency special tokens discussed in LLM studies, yet extreme outliers also emerge in Vision Transformers (ViT) [20, 31]. Thus, we argue that the root cause of extreme outliers is structural, while data properties may still modulate their exact locations or magnitudes.

When the cause is “mechanical”, so should be the solution for extreme outliers. We propose TWEO (Transformers Without Extreme Outliers), a loss function that directly penalizes extreme activation magnitudes. TWEO has the following core advantages:

- **Universal:** It is data-independent, which directly suppresses extreme magnitudes regardless of their origin or functionality. Thus, TWEO is naturally cross-modal.
- **Flexible:** It can be plugged in to any Transformer architecture or variants.
- **Effective:** It eliminates loss explosion and ensures smooth and effective FP8 training.
- **Simple:** It can be used *out of the box* (i.e., “FP8 training for dummies”), without resorting to engineering tricks. Another benefit is that TWEO-trained models can be effectively quantized by the simplest method (i.e., “quantization for dummies”).

2. Related Work

Low-bit computation for Transformer models (FP8 training or quantization) are seriously hindered by extreme outliers. We briefly outline the solutions in the literature.

Stabilize FP8 Training. When extreme outliers exceed the FP8 range, catastrophic collapse happens and disables FP8 training [12, 18, 28]. Current solutions include

1. Mixed-precision engineering [25, 28, 40, 46]. These methods keep sensitive modules (e.g., embeddings, normalization, MoE gates) in BF16, but entail high engineering costs and sacrifice FP8’s efficiency benefits.
2. Data-dependent architectural modifications. They introduce register tokens like ViT-R [8, 10, 35] to absorb outliers. But, we will show that extreme outliers are in fact data-independent.
3. Data-independent invasive modifications. These methods propose invasive architectural changes such as modifying activation functions (Smooth-SwiGLU [15]), attention (Clipped Softmax [4], OutEffHop [23]), or com-

pletely redesigning the architecture (FOG [22]).

These solutions are either costly, based on inappropriate assumptions, or lack generality. The community needs a non-invasive, plug-and-play, universal solution for stable FP8 training without heavy engineering. We propose TWEO to approach this goal.

Simple and effective PTQ. In Post-Training Quantization (PTQ, [6, 16]), extreme outliers drastically stretch the quantization range and lead to poor accuracy [34, 53]. To address this challenge, existing methods (e.g., SmoothQuant, OS+, AWQ, and QuaRot) have to transfer the difficulty or remove outliers through equivalent transformations or rotations [2, 50]. They concede that activations (especially the residual stream [41]) are difficult to quantize due to outliers, thus use mathematical transformations to transfer the quantization difficulty from the hard activations to the relatively easy weights. SmoothQuant [51] introduces offline scaling factors, OS+ [50] uses channel-wise shifting and scaling, while AWQ [27] scales activations by protecting salient weights. SpinQuant [30] and QuaRot [2] introduce rotations to reduce or eliminate the impact of outliers, while compensation methods attempt to fix or compensate for errors post-quantization [17]. Related work also explores activation regularization to mitigate outlier channels during training for low-bit quantization [33].

Although these methods have been effective to some extent in practice, they increase algorithmic and inference complexity. More importantly they are *forced to keep the residual stream in high precision* (cf. Sec. 4.4 and Appendix D), which introduces massive type-conversion operations and memory overhead and hinders acceleration in real-world scenarios.

These drawbacks stem from the fact that they assume extreme outliers are unavoidable. Our TWEO eliminates extreme outliers, hence makes the simplest per-tensor static quantization viable (*including the residual stream*), and renders complex methods (like SmoothQuant) unnecessary.

3. Method

We start by examining existing assumptions on cause of extreme outliers, presenting our assumption, and then presenting the TWEO solution.

3.1. Data-dependent Assumptions Are Problematic

Existing work often attributes extreme outliers to properties of the data, e.g., token frequency or special tokens [4, 7, 8, 35]. Following [42], we propose a *Contradiction Stethoscope* to test whether these factors are necessary for the emergence of extreme outliers.

Assumption 1 (The Contradiction Stethoscope)

Suppose the input data are replaced by random samples that deliberately remove the special structures invoked

by data-dependent assumptions. If extreme outliers still persist at a similar scale, then these data properties are not necessary for producing the outliers.

As shown in Figure 1, even when the inputs are replaced by random samples, the pre-trained model still produces extreme outliers. For LLMs, we sample tokens uniformly from the full vocabulary using `torch.randint`; for ViTs, we sample pixel values uniformly in $[0, 255]$ using `np.random.randint` (Appendix A). These maximum-entropy inputs remove semantic structure, long-tail token frequency, and spatial correlations. In contrast, when real input data are fed into a *randomly initialized* model, no extreme outliers appear. Hence, the evidence indicates that the root cause is associated with the *trained weights*, rather than being induced by specific input semantics.

For each layer, the *Contradiction Stethoscope* reports three simple statistics over the batch: Avg Max, Avg Min, and Mean-Abs. These metrics provide a compact layer-wise summary of whether extreme positive or negative activations persist under different model/input combinations.

Hence, our data-independent assumption is

Assumption 2 (Data-independent Root Cause) *Extreme outliers are primarily a mechanical artifact resulting from specific weight structures acquired during training, rather than being created by particular input semantics or token/image statistics.*

We provide more exhaustive studies spanning both language and vision in Appendix A, where we further examine the roles of model weights, token embeddings, and input distributions. These results show that outliers emerge with high consistency across real data and random inputs once the trained weights are present.

3.2. Colinearity Causes Extreme Outliers

Recent studies [15, 20, 31] indicate that extreme outliers primarily originate in the MLP (Multi-Layer Perceptron) layers of the Transformer blocks. Hence, we conducted an empirical analysis of both classic MLPs with GELU activation [21] and modern MLPs with GLU variants.

Consider a simplified classic MLP layer, whose computation can be approximately represented as $y = BAx$ by temporarily ignoring the activation function. We analyze the conditions under which the k -th element y_k of the output $y \in \mathbb{R}^{d_1}$ becomes an outlier. y_k can be expressed as the product of the k -th row of the matrix $B \in \mathbb{R}^{d_1 \times d_2}$ (denoted as $w^T \in \mathbb{R}^{1 \times d_2}$) and Ax .

Performing Singular Value Decomposition (SVD) on the matrix $A \in \mathbb{R}^{d_2 \times d_1}$ (up-projection) as $A = \sum_{i=1}^{d_1} s_i u_i v_i^T$, y_k can be rewritten as

$$y_k = w^T Ax = \sum_{i=1}^{d_1} s_i (w^T u_i) (v_i^T x), \quad (1)$$

which suggests a potential cause of outliers: when the row vector w of matrix B (down-projection) becomes *collinear* with a left singular vector u_i of matrix A , the product $(w^T u_i)$ becomes very large. If, at the same time, the input x is highly aligned with the corresponding right singular vector v_i of A (which causes $v_i^T x$ to be large), y_k will become an extreme outlier due to the multiplication of these two large values.

We verified this hypothesis on a ViT-B model, ViT-B-21k which uses a classic GELU MLP. We located the first MLP layer to exhibit extreme outliers, with a value of 880. We then performed SVD analysis on the weight matrices A and B of this layer, i.e., $y = B \cdot \text{gelu}(A \cdot x)$. Analyzing only the dominant singular vector ($i = 1$), we found that $s_1 = 18.189$, $w^T u_1 = -9.1847$ (which showed high weight and singular vector alignment), and $v_1^T x = -5.6165$ (which showed high input and singular vector alignment). Based on Eq. 1 (our analysis which ignores GELU), the simulated value is 938.28. When we further incorporated the GELU activation function, the simulated value became 884.29.

Note that the simulated value (884.29) is almost identical to the true outlier (880, relative error $< 0.5\%$). Beyond this simple example, we find that whenever extreme outliers happened, Eq. 1 almost always simulated them accurately. This fact strongly suggests that outliers are a mechanical artifact resulting from structural alignment between the weight matrices (A, B) and the input (x).¹

In architectures with the GLU activation function (e.g., in Llama[14, 47, 48], Qwen[3, 52]), the effect of Eq. 1 is further amplified by the (anti-)colinearity between the corresponding row vectors of the gate_proj and up_proj weight matrices [15], which leads to even larger extreme outliers.

Combining the empirical evidence for Assumption 2 and the theoretical analysis in Eq. 1, our main hypothesis is: **extreme outliers are a structural artifact**. That is, they are not (or not solely) a direct representation of input data semantics (like special tokens), but are the result of mechanical amplification caused by specific properties that the network’s weight matrices gradually develop during training dynamics (e.g., singular vector alignment or colinearity).

Then, we argue that to remove extreme outliers, we should not rely on modifying specific architectures (like the MLP or Attention), but design a method that directly suppresses this amplification effect during training mechanically.

¹When Eq. 1 predicts an extreme outlier, the real value may not be an extreme outlier because the activation function is ignored in Eq. 1. In practice, these false positives are typically values that are amplified by the linear approximation but then dampened by the true activation function (e.g., GELU). False positives, however, are not important, because our Assumption 2 focuses only on the root cause of extreme outliers.

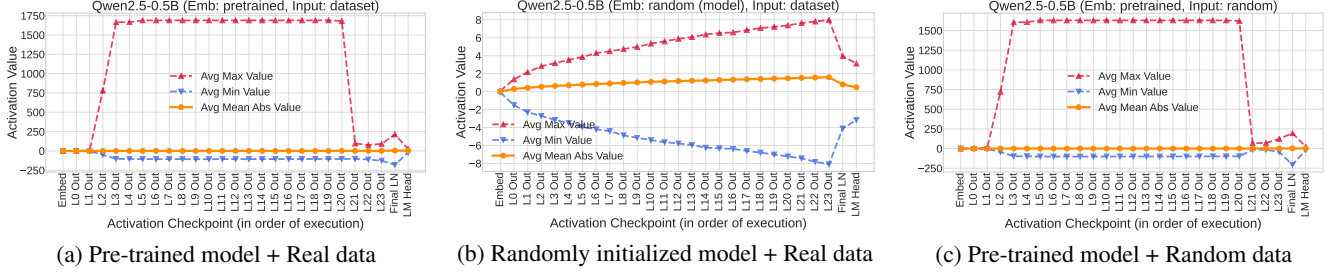


Figure 1. Empirical evidence for data-dependent vs. data-independent assumptions. (a) Running real data through a pre-trained Qwen2.5-0.5B LLM produces extreme outliers (>1000) in most layers. (b) But, real data through a *randomly initialized* Qwen2.5-0.5B yields minimal activation magnitudes (<10). (c) And, changing the input to *random token inputs*, the same pre-trained Qwen2.5-0.5B **still produces** extreme outliers.

3.3. The Proposed TWEO Loss

We then propose the TWEO (Transformers Without Extreme Outliers) loss $\mathcal{L}_{\text{TWEO}}$. It is a regularization term added to the main task loss, $\mathcal{L}_{\text{task}}$:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{task}} + \lambda(t)\mathcal{L}_{\text{TWEO}}, \quad (2)$$

where the hyperparameter $\lambda(t)$ varies with training step t .

Unlike previous studies that might focus on internal MLP activations, our method intervenes at a more macroscopic scale. We monitor the final output activation $A^{(l)}$ of a Transformer block across all blocks $l \in \{1, \dots, L\}$, i.e., the y after $y = x + \text{MLP}(\text{LN}(x))$.

$\mathcal{L}_{\text{TWEO}}$ is defined as the average scaled L_p loss computed over the activations $A^{(l)}$ from all L Transformer blocks:

$$\mathcal{L}_{\text{TWEO}} = \frac{1}{L} \sum_{l=1}^L \mathbb{E} \left[\left(\frac{|A^{(l)}|}{\tau + \epsilon} \right)^p \right], \quad (3)$$

where $|\cdot|$ computes the absolute value, $A^{(l)}$ is the activation tensor of the l -th Transformer block, $\mathbb{E}[\cdot]$ denotes the mean (average) taken over the batch, sequence length, and hidden dimensions, $\tau > 0$ is a magnitude scaling factor or soft threshold, p is the power of the penalty, and ϵ is a small constant (e.g., $1e-6$) to ensure numerical stability.

The core of this design is to create a non-linear regularization mechanism that perfectly matches our goal of *preserving normal values while suppressing extreme values*. This is achieved through the synergy of two key parameters: the magnitude scaling factor τ and the penalty power p (where $\tau \in [2, 5]$, $p = 4$ in our design).

Unlike the Hinge loss $\max(0, |A| - \tau)$, which uses a hard threshold, our τ acts as a soft threshold or scaling factor. It does *not* introduce discontinuous gradients, which helps training stability, and it defines the target scale for activation magnitudes.

The penalty power $p = 4$ is key to achieving this non-linear penalty. The interaction of τ and p is as follows.

- Tolerance for normal values (when $|A| < \tau$). For activations much smaller than τ (e.g., $|A| = 0.5\tau$), the power

$p = 4$ renders its penalty term very small ($(0.5)^4 = 0.0625$). This allows the model to freely retain low-magnitude normal activations.

- Setting the target scale (when $|A| = \tau$). When the activation equals the desired scale, the penalty for this term is a medium one: $(1)^p = 1$.
- Suppression of extreme outliers (when $|A| > \tau$). For activations exceeding τ (e.g., $|A| = 10\tau \in [20, 50]$), the penalty term is $(10)^4 = 10000$, which is a massive cost that will aggressively suppress extreme outliers.

TWEO is then highly effective at suppressing the extreme tail of the activation value’s distribution, which is the root cause of FP8 overflows, while not affecting the training of normal activation values.

TWEO is simple and flexible, which allows it to be applied to various Transformer architectures. Because it acts directly on the physical magnitude of activations, rather than relying on specific task semantics (like token semantics in language or image structures in vision), it can be seamlessly extended to the training of vision, language, and even generative models (as shown in Appendix G).

4. Experiments

We designed a series of experiments aimed at answering the following key questions: (1) Does TWEO possess cross-modal (vision, language) universality? (2) Can TWEO solve the loss explosion and instability issues in native FP8 pre-training? (3) What positive downstream impacts (e.g., on quantization) do TWEO have?

4.1. Experimental Setup

Models and Datasets Our experiments cover two mainstream modalities:

- **Vision:** We trained Swin Transformers (Swin-T, Swin-S) [29] on the ImageNet-1K [11] dataset.
- **Language:** We pre-trained a series of GPT-2 models [39] from scratch on the OpenWebText [19] dataset, with parameter scales ranging from 124M (GPT2) to 1.5B (GPT2-XL), as well as scaling to 3B (GPT2-3B) and

Table 1. Results on Swin and ViT models on ImageNet. #Params: parameter counts. Mod.: whether the network architecture has been modified. Peak: largest outlier value during training. Final: largest outlier value after training was finished.

Method	Params	Mod.	Top-1 (%)	Peak	Final
Swin-T [29] (Baseline)	28M	-	81.2	1556	534
Swin-T attn bias [1]	28M	Yes	81.1	478	135
Swin-T softmax+1 [4, 24]	28M	Yes	81.4	811	143
Swin-T gated [37]	28M	Yes	80.7	976	126
Swin-T TWEO	28M	No	81.4	22	15
Swin-S [29] (Baseline)	50M	-	82.7	6402	1758
Swin-S TWEO	50M	No	82.8	35	18
Swin-B [29] (Baseline)	88M	-	83.5	4824	876
Swin-B TWEO	88M	No	83.4	37	16
ViT-S [13] (Baseline)	22M	-	79.8	328	174
ViT-S TWEO	22M	No	79.6	36	19
ViT-B [13] (Baseline)	87M	-	81.3	1579	106
ViT-B TWEO	87M	No	81.3	38	16

7B (GPT2-7B) using configurations similar to GPT-3 [5]. Configuration details are available in Appendix C.

We compared against BF16 training, standard native FP8 training (which led to training collapse), and outlier suppression methods in the literature [4, 10].

For TWEO, we applied it to the final output activations of all Transformer blocks. We found the settings for τ (scaling factor) and p (penalty power) are robust across all tasks, with $\tau = 3$ and $p = 4$ in *all our experiments* if not otherwise specified. By default, we use a constant $\lambda(t) = 0.01$ throughout training; cosine annealing is only an optional variant for mildly reducing the regularization strength near convergence.

4.2. Results on Vision Tasks

We trained Swin Transformers and ViT on ImageNet and compared against invasive methods that require architectural changes. As shown in Table 1, we can observe

- Existing methods require modifying the model’s attention or activation modules. TWEO achieves similar accuracy to the baselines *without any structural modifications*.
- TWEO eliminates extreme outliers while existing methods cannot. For Swin-S, it reduced the training peak magnitude from 6402 to 35.2 (99.4% reduction), while for Swin-T from 1556 to 21.6 (98.6% reduction).

This set of experiments confirms that the extreme outlier issue also exists in vision models and that TWEO, as a data-independent method, can effectively solve it.

4.3. LLMs: Stabilizing Full FP8 Pre-training

This is the core application scenario for TWEO. To clearly demonstrate TWEO’s simplicity and efficiency, we com-

Table 2. FP8 pre-training. PPL: Validation perplexity. ↓: lower is better. Peak: largest outlier during training.

Model	Params	PPL (FP8)	PPL (BF16)	Peak
GPT2 (baseline)	124M	169.81	20.04	823
Attn bias * [1]	124M	24.93	19.79	163
Attn ca-softmax * [1]	124M	23.81	19.34	135
Baseline+TWEO	124M	19.26	18.68	17
GPT2-medium	350M	127.34	16.77 #	24563
+TWEO	350M	15.64	15.18	19
GPT2-large	774M	178.66	14.78 #	32361
+TWEO	774M	13.89	13.79	18
GPT2-xl	1.5B	93.28	13.84	32889
+TWEO	1.5B	12.58	12.39	19
GPT2 3B	3B	76.85	-	-
+TWEO	3B	12.24	-	18
GPT2 7B	7B	67.12	-	-
+TWEO	7B	12.02	-	20

* Invasive methods.

Attn bias: Adds learnable biases k' and v' to the Key and Value matrices. Attn ca-softmax: Adds a learnable, context-dependent scaling factor.

Training collapsed. We report the best checkpoint before collapsing.

pare it with the FP8 training solution of the well-known DeepSeek-V3 [28] model.

DeepSeek-V3 stated that low-precision training is “limited by the presence of outliers in activations, weights, and gradients” [28]; hence, it is a highly complex, deeply customized *mixed-precision* engineering effort:

- Bypassing Outliers: It is forced to keep all outlier-prone areas, such as the embedding module, language model head, and normalization operators in BF16 or FP32.
- Isolating Outliers: For modules using FP8, it relies on **fine-grained** ‘tile-wise’ or ‘block-wise’ quantization strategies, calculating separate scaling factors for different blocks of a tensor to isolate local outliers.

This approach requires many customized operators (e.g., online quantization and CUDA Core accumulation [28]), making it difficult to transfer and popularize.

In contrast, the core advantage of our TWEO is that it *eliminates outliers from the root cause* (cf. Figures 5 and 4). This obviates the need for complex engineering workarounds to adapt to or bypass outliers. Hence, we are able to adopt a *very aggressive* and *minimalist* training strategy. We placed all Linear layers and LayerNorm components (including the final language model head) under the FP8 automatic mixed precision (autocast) context for computation, a setup that is unacceptable in frameworks like DeepSeek-V3.

Our FP8 training setup aims for maximum simplicity. We utilize the NVIDIA Transformer Engine [22], em-




Figure 2. Activation magnitudes in GPT-2 Medium BF16 training.

ploying the *coarse-grained* DelayedScaling strategy, which has the lowest hardware overhead but **zero tolerance** for outliers. It uses *one single scaling factor for the entire tensor* (true per-tensor scaling) and the `Format.HYBRID` (E4M3/E5M2) format.

Its aggressive nature is particularly evident in our configuration of an *extremely short* `amax_history_len=16`. This is a high-risk configuration: it assumes a stable activation distribution (since ‘per-tensor’ scaling cannot isolate local spikes) and is thus extremely prone to training collapse when an outlier emerges.

As Figures 2 and 3 show, this high-risk setup perfectly validates our hypothesis. The standard model (Baseline) suffered a catastrophic loss explosion under this stringent configuration. In contrast, TWEO kept a smooth and stable convergence trajectory identical to the BF16 baseline, even under this simplest, most aggressive per-tensor FP8 setting.

Table 2 summarizes these experiments. Standard FP8 training (Baseline) suffered catastrophic collapse across all model scales from 124M to 7B. Its validation perplexity (PPL) is extremely high (e.g., 93.28 for GPT2-XL, 67.12 for GPT2 7B). In contrast, TWEO (‘+TWEO’ in Table 2) not only successfully completed training at all scales, but their FP8 PPL matched or even surpassed those of the BF16 baseline (e.g., 12.58 for GPT2-XL, 12.02 for GPT2 7B). The successful suppression of extreme outliers is the key: as shown in the last column, the baseline models’ peak magnitudes reached tens of thousands, whereas TWEO consistently controlled the peak magnitudes of all models to ≤ 20 .

Figure 4 visualizes TWEO’s outlier suppression effect from a distributional perspective. The baseline model (‘Default’) had many modes and was heavy-tailed, whose extreme values (e.g., ± 30000 in GPT2-XL) caused the distribution to be extremely stretched. In contrast, TWEO’s activation distributions are unimodal, with means near 0 and small standard deviations, and the extreme outliers are eliminated.

Figure 5 intuitively demonstrates TWEO’s mechanical suppression effect. The baseline model’s training peak magnitude reached 30000, while TWEO’s activation values across *all layers* were smoothly clamped below 20.

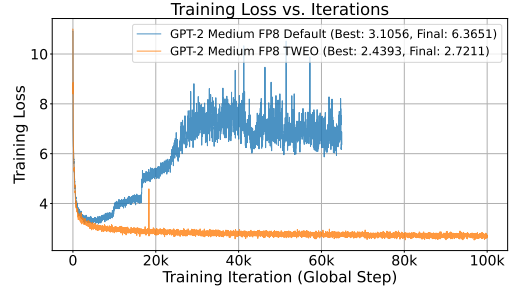


Figure 3. Activation magnitudes in GPT-2 Medium FP8 training.

4.4. Post Training: A New Quantization Landscape

The value of TWEO lies not only in enabling stable low-bit training but also in facilitating post-training quantization. A consensus is that extreme activation outliers are the major barrier to quantization [26, 53]. This issue has led the field to *abandon the use of per-tensor static quantization for activations*, because per-tensor quantization is the most hardware-friendly and fastest inference solution, but it is also the most sensitive to outliers. We have been forced to use per-token or other more expensive quantization schemes.

Now that TWEO has removed this barrier (extreme outliers), it must enable simple and effective post-training quantization. To verify this, we applied the strictest (most outlier-sensitive) and simplest (most hardware-friendly) *AbsMax static quantization* to both baseline and TWEO-trained models. AbsMax static quantization is simple, symmetric, and hardware-friendly (zero-point-free). We use the ‘‘WbAa’’ notation (e.g., ‘‘W8A8’’) to denote weights quantized to b bits and activations to a bits.

For a b -bit symmetric quantization, we first define the maximum quantization integer $Q_b = 2^{b-1} - 1$. For a floating-point tensor X , the quantization process is:

$$s = \max(|X|) + \epsilon, X_q = \text{round}\left(\frac{X}{s} \cdot Q_b\right), X_{\text{deq}} = \frac{X_q}{Q_b} \cdot s \quad (4)$$

The speed bottleneck lies in the granularity of the scaling factor s calculation, as:

- Per-tensor (T): (Used for W and A) A single scaling factor s for the *entire* tensor.
- Per-channel (C): (Used for W only) Independent s_c for *each* output channel.
- Per-token (K): (Used for A only) Independent s_t for *each token vector*. Highest overhead.

We compare combinations of these granularities to comprehensively evaluate TWEO’s impact on PTQ.

Results in Table 3 clearly demonstrate a paradigm shift:

- As expected, the Default (Baseline) models completely collapsed in performance whenever per-tensor activation quantization, A(T) or W(T)+A(T), was involved. For ex-

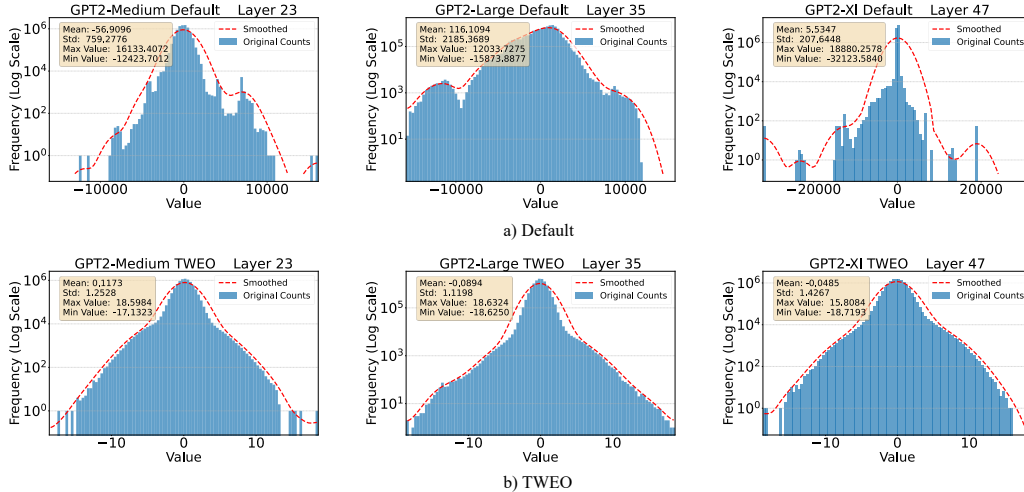


Figure 4. Activation value distributions for various GPT-2 model sizes, with and without TWEO.

Table 3. Perplexity (PPL) comparison of Baseline vs. TWEO models under different 8-bit AbsMax PTQ strategies on OpenWebText.

Model (Params)	Method	BF16 PPL	Act. Only 8-bit (A8)		Weight Only 8-bit (W8)		Full 8-bit Quant (W8A8)			
			per-tensor	per-token	per-tensor	per-channel	W8(T)A8(T)	W8(T)A8(K)	W8(C)A8(T)	W8(C)A8(K)
GPT-2 (124M)	Default	20.04	82.94	21.19	20.71	20.32	86.60	22.02	82.94	21.49
	TWEO	18.83	20.21	19.00	19.32	19.11	20.82	19.51	20.54	19.00
GPT-2 Medium (350M)	Default	16.77	1451.40	19.76	17.35	17.17	1491.11	20.24	1456.97	19.95
	TWEO	15.18	16.28	15.34	15.38	15.25	16.50	15.53	16.32	15.40
GPT-2 Large (0.8B)	Default	14.92	43204.39	23.20	2228.19	2228.19	43720.76	24.14	45821.84	23.35
	TWEO	13.79	18.09	13.98	13.98	13.89	18.88	14.17	18.39	14.07
GPT-2 XL (1.5B)	Default	13.84	1799.44	22.42	14.27	14.08	1872.83	22.01	1799.44	23.35
	TWEO	12.39	12.86	12.43	12.61	12.58	13.09	12.66	13.06	12.63

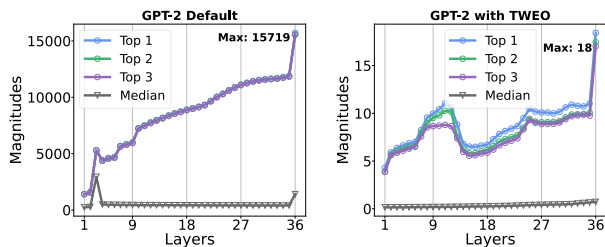


Figure 5. Comparison of activation magnitudes during GPT-2 Large (774M) training. Top 1 refers to the largest activation value in each layer. Top 2, Top 3 refer to the second and third largest. Median is the median of all activations.

ample, the PPL of GPT-2 Medium spiked from 16.77 to 1451.40 or 1491.11.

- But, activations are no longer the bottleneck in TWEO. On GPT-2 XL trained by TWEO, the per-token A8(K) quantization PPL (12.43) is even better than its per-channel W8(C) PPL (12.58). This observation subverts the traditional wisdom “activations are always harder to

quantize than weights.”

- The critical finding comes from the W(T)+A(T) (*full per-tensor*) strategy. For GPT-2 Medium, TWEO achieved a PPL of 16.50. This result is significant because a) It stands in stark contrast to the baseline model’s PPL (1491.11), and b) It is even better than the Default (Baseline)’s BF16 PPL (16.77). This result shows that by eliminating outliers with TWEO, we *may not need* the complex difficulty transfer techniques of SmoothQuant or SpinQuant. Instead, we just need to directly employ the fastest, simplest-to-implement per-tensor static quantization scheme while achieving low performance loss.

Table 3 proves that TWEO can make simple AbsMax quantization viable, but a more profound problem lies in the quantization of the **residual stream**. In the Transformer architecture, activations are passed between layers via a residual stream $y = x + f(x)$. Extreme outliers are passed and accumulated in the residual stream layer by layer. Consequently, difficulty transfer methods like SmoothQuant are forced to operate *only* inside $f(x)$ while being compelled to leave x itself (the residual stream) in high precision

Table 4. Perplexity of GPT-2 XL. SmoothQuant collapsed when the residual was quantized. More details are provided in Appendix D. Res.?: whether the residual stream was quantized or not.

Model PPL (BF16)	Method	Res.?	W8(C)A8(T) / W8(C)A8(K)
GPT-2 XL Default (PPL: 13.84)	SmoothQuant	No	14.81 / 14.01
	SmoothQuant	Yes	1876.70 / 21.93
GPT-2 XL TWEO (PPL: 12.39)	AbsMax	Yes	13.06 / 12.63
	SmoothQuant	Yes	12.89 / 12.51

(BF16/FP32). As shown in Table 4, quantizing x led to catastrophic collapse. This compromise (not quantizing the residual) results in frequent BF16 \leftrightarrow int8 type conversions and introduces significant memory and latency overhead, severely hampering inference speed.

By eliminating extreme outliers, TWEO makes the quantization of the residual stream x possible *for the first time*. Results in Table 4 show that

- When SmoothQuant was forced to quantize the residual stream, its W8(C)A8(T) PPL on the Default model catastrophically collapsed from 14.81 to 1876.70.
- Conversely, the TWEO model, even when using AbsMax and quantizing the residual stream, maintained a PPL of 13.06 / 12.63, which is even better than all baselines.
- On the TWEO pre-trained model, applying the complicated SmoothQuant method yielded negligible gains compared to simple AbsMax. This shows that TWEO has eased difficulties in quantization of *both* activations and weights, making complex difficulty transfer methods unnecessary.

That is, our work introduces a new paradigm for post-training quantization, which makes full-model quantization (*including the residual stream*) possible and paves the way for unleashing the full potential of low-bit inference. We also observe similar PTQ improvements on vision transformers under the strict W(C)A(T) setting; detailed ViT and Swin results are deferred to Appendix F.

Implications for the Quantization Field. That is, TWEO has the potential to open up a new, outlier-free research direction for the field of quantization.

For academic research, the future focus may shift from “how to design complex algorithms to bypass outliers” to “with an outlier-free model, how low-bit can we go?” TWEO simplifies the quantization problem itself, allowing researchers to focus on exploring more aggressive low-bit (e.g., W4A8, W4A4) compression schemes for this clean (i.e., outlier-free) class of models.

For industry applications, a scheme using per-tensor static quantization for both activations and weights such as W8(T)A8(T) is unrivaled in its appeal. It requires no complex re-scaling or online rotations, supports maximum concurrency, and enjoys minimum latency, power consump-

tion, and memory footprint. TWEO proves that this simplest, most efficient path is in fact viable, and the substantially reduced costs can drive on-device deployment of more artificial intelligence products.

Furthermore, TWEO opens new paths for hardware-software co-design. Previous ASICs (Application-Specific Integrated Circuits) had to build costly control logic for complex per-token/per-channel quantization schemes (to circumvent the outlier challenges [51]). This not only increases chip area and power consumption but also restricts their maximum clock frequency [45, 49].

TWEO’s potential contribution in this aspect is making the simplest, most hardware-friendly *static per-tensor quantization* viable again (as shown in Table 3). This is a significant benefit for next-generation specialized AI accelerators: designers can remove complex quantization logic and instead design leaner, more specialized computation cores. This simplified design has the potential to translate into higher clock frequencies at lower power, enabling a leap in energy efficiency [9, 45].

5. Conclusions

In this paper, we confronted one of the core bottlenecks hindering next-generation AI hardware from reaching its full potential: extreme activation outliers. Instead of treating them as symptoms to be bypassed, we hypothesized that their root cause is *data-independent, mechanical artifacts* of the training dynamics. Then we designed TWEO, a plug-and-play, non-invasive loss function, which cures the model by proactively suppressing the formation of extreme magnitudes at every step of training.

Our experiments on both language and vision up to 7B models demonstrated the impact of this approach:

- *In FP8 Training:* It transforms catastrophic, unreliable full FP8 training into a stable, efficient (+36% throughput) process with performance on par with BF16.
- *In Post-Training Quantization:* It allows the field to move beyond its dependency on complex difficulty transfer algorithms (like SmoothQuant). We demonstrated that on an outlier-free model produced by TWEO, the most hardware-friendly, simplest *per-tensor static quantization*—including *full residual stream* quantization—becomes viable and effective for the first time.

TWEO also moves low-bit computation from expert customization to out-of-the-box technology.

Due to resource constraints, we had not verified TWEO on the largest models (e.g., 700B), which will be a future work when compute is available. While TWEO models were trained from scratch in this paper, we are also interested in fine-tuning existing models to remove their extreme outliers. And, we are interested in further FP4 learning and inference.

Acknowledgments

This work was partly supported by the National Natural Science Foundation of China under Grant 62276123, by the Zhongguancun Academy Project No. XTS0066, and by the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (No. JYB2025XDXM118).

J.W. designed the strategy to eliminate extreme outliers via finding its root cause, and then enables easy FP8 training and quantization. J.W., with the help of G.L., proposed the *Contradiction Stethoscope*. G.L., with the help of J.W., captured the data-independent nature of outliers, proposed the TWEQ loss, and successfully implemented it. J.S., N.T., and X.L. provided essential support in carrying out the experiments and analyses. All authors contributed to the final manuscript.

References

- [1] Yongqi An, Xu Zhao, Tao Yu, Ming Tang, and Jinqiao Wang. Systematic outliers in large language models. *arXiv preprint arXiv:2502.06415*, 2025. 1, 5
- [2] Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefer, and James Hensman. Quarot: Outlier-free 4-bit inference in rotated llms. *Advances in Neural Information Processing Systems*, 37:100213–100240, 2024. 2
- [3] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 3
- [4] Yelysei Bondarenko, Markus Nagel, and Tijmen Blankevoort. Quantizable transformers: Removing outliers by helping attention heads do nothing. *Advances in Neural Information Processing Systems*, 36:75067–75096, 2023. 1, 2, 5
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 5
- [6] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13169–13178, 2020. 2
- [7] Kexin Chen, Dongxia Wang, Yi Liu, Haonan Zhang, and Wenhai Wang. Sticking to the mean: Detecting sticky tokens in text embedding models. *arXiv preprint arXiv:2507.18171*, 2025. 1, 2
- [8] Mengzhao Chen, Yi Liu, Jiahao Wang, Yi Bin, Wenqi Shao, and Ping Luo. Prefixquant: Eliminating outliers by prefixed tokens for large language models quantization. *arXiv preprint arXiv:2410.05265*, 2024. 1, 2
- [9] NVIDIA Corporation. Tensorrt: High-performance deep learning inference library. <https://github.com/NVIDIA/TensorRT>, 2023. Accessed: 2023-10-01. 8
- [10] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023. 1, 2, 5
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4
- [12] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in neural information processing systems*, 35:30318–30332, 2022. 1, 2
- [13] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5
- [14] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv e-prints*, pages arXiv-2407, 2024. 3
- [15] Maxim Fishman, Brian Chmiel, Ron Banner, and Daniel Soudry. Scaling fp8 training to trillion-token llms. *arXiv preprint arXiv:2409.12517*, 2024. 2, 3
- [16] Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022. 2
- [17] Minghao Fu, Hao Yu, Jie Shao, Junjie Zhou, Ke Zhu, and Jianxin Wu. Quantization without tears. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 4462–4472, 2025. 2
- [18] Jorge Gallego-Feliciano, S Aaron McClendon, Juan Morinelli, Stavros Zervoudakis, and Antonios Saravanos. Hidden dynamics of massive activations in transformer training. *arXiv preprint arXiv:2508.03616*, 2025. 1, 2
- [19] Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. The Pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020. 4
- [20] Bobby He, Lorenzo Noci, Daniele Paliotta, Imanol Schlag, and Thomas Hofmann. Understanding and minimising outlier features in transformer training. *Advances in Neural Information Processing Systems*, 37:83786–83846, 2024. 2, 3
- [21] D Hendrycks. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 3
- [22] Alejandro Hernández-Cano, Dhia Garbaya, Imanol Schlag, and Martin Jaggi. Towards fully fp8 gemm llm training at scale. *arXiv preprint arXiv:2505.20524*, 2025. 1, 2, 5
- [23] Jerry Yao-Chieh Hu, Pei-Hsuan Chang, Haozheng Luo, Hong-Yu Chen, Weijian Li, Wei-Po Wang, and Han Liu. Outlier-efficient hopfield layers for large transformer-based models. In *Proceedings of the 41st International Conference on Machine Learning*, pages 19123–19152, 2024. 2

- [24] Prannay Kaul, Chengcheng Ma, Ismail Elezi, and Jiankang Deng. From attention to activation: Unravelling the enigmas of large language models. *arXiv preprint arXiv:2410.17174*, 2024. 5
- [25] Ang Li, Ben Liu, Binbin Hu, Bing Li, Bingwei Zeng, Borui Ye, Caizhi Tang, Changxin Tian, Chao Huang, Chao Zhang, et al. Every activation boosted: Scaling general reasoner to 1 trillion open language foundation. *arXiv preprint arXiv:2510.22115*, 2025. 2
- [26] Guang Liang, Xinyao Liu, and Jianxin Wu. Gplq: A general, practical, and lightning qat method for vision transformers. *arXiv preprint arXiv:2506.11784*, 2025. 6
- [27] 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
- [28] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024. 1, 2, 5
- [29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 4, 5
- [30] 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. *arXiv preprint arXiv:2405.16406*, 2024. 2
- [31] Yuexiao Ma, Huixia Li, Xiawu Zheng, Feng Ling, Xuefeng Xiao, Rui Wang, Shilei Wen, Fei Chao, and Rongrong Ji. Outlier-aware slicing for post-training quantization in vision transformer. In *Forty-first International Conference on Machine Learning*, 2024. 2, 3
- [32] Paulius Micikevicius, Dusan Stosic, Neil Burgess, Marius Cornea, Pradeep Dubey, Richard Grisenthwaite, Sangwon Ha, Alexander Heinecke, Patrick Judd, John Kamalu, et al. Fp8 formats for deep learning. *arXiv preprint arXiv:2209.05433*, 2022. 1
- [33] Aniruddha Nrusimha, Mayank Mishra, Naigang Wang, Dan Alistarh, Rameswar Panda, and Yoon Kim. Mitigating the impact of outlier channels for language model quantization with activation regularization. *arXiv preprint arXiv:2404.03605*, 2024. 2
- [34] Jungwoo Park, Taewhoo Lee, Chanwoong Yoon, Hyeon Hwang, and Jaewoo Kang. Outlier-safe pre-training for robust 4-bit quantization of large language models. *arXiv preprint arXiv:2506.19697*, 2025. 1, 2
- [35] Giovanni Puccetti, Anna Rogers, Aleksandr Drozd, and Felice Dell’Orletta. Outliers dimensions that disrupt transformers are driven by frequency. *arXiv preprint arXiv:2205.11380*, 2022. 1, 2
- [36] Haiquan Qiu and Quanming Yao. Why low-precision transformer training fails: An analysis on flash attention. *arXiv preprint arXiv:2510.04212*, 2025. 1
- [37] Zihan Qiu, Zekun Wang, Bo Zheng, Zeyu Huang, Kaiyue Wen, Songlin Yang, Rui Men, Le Yu, Fei Huang, Suozhi Huang, et al. Gated attention for large language models: Non-linearity, sparsity, and attention-sink-free. *arXiv preprint arXiv:2505.06708*, 2025. 5
- [38] Zihan Qiu, Zeyu Huang, Kaiyue Wen, Peng Jin, Bo Zheng, Yuxin Zhou, Haofeng Huang, Zekun Wang, Xiao Li, Huaqing Zhang, et al. A unified view of attention and residual sinks: Outlier-driven rescaling is essential for transformer training. *arXiv preprint arXiv:2601.22966*, 2026. 1
- [39] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 4
- [40] Bitan Darvish Rouhani, Ritchie Zhao, Ankit More, Mathew Hall, Alireza Khodamoradi, Summer Deng, Dhruv Choudhary, Marius Cornea, Eric Dellinger, Kristof Denolf, et al. Microscaling data formats for deep learning. *arXiv preprint arXiv:2310.10537*, 2023. 2
- [41] Utkarsh Saxena, Sayeh Sharify, Kaushik Roy, and Xin Wang. Resq: Mixed-precision quantization of large language models with low-rank residuals. *arXiv preprint arXiv:2412.14363*, 2024. 2
- [42] Jie Shao and Jianxin Wu. Who reasons in the large language models? *arXiv preprint arXiv:2505.20993*, 2025. 2
- [43] Mingjie Sun, Xinlei Chen, J Zico Kolter, and Zhuang Liu. Massive activations in large language models. *arXiv preprint arXiv:2402.17762*, 2024. 1
- [44] Shangwen Sun, Alfredo Canziani, Yann LeCun, and Jiachen Zhu. The spike, the sparse and the sink: Anatomy of massive activations and attention sinks. *arXiv preprint arXiv:2603.05498*, 2026. 1
- [45] Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017. 8
- [46] Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv preprint arXiv:2507.20534*, 2025. 2
- [47] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 3
- [48] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 3
- [49] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8612–8620, 2019. 8
- [50] Xiuying Wei, Yunchen Zhang, Yuhang Li, Xiangguo Zhang, Ruihao Gong, Jinyang Guo, and Xianglong Liu. Outlier suppression+: Accurate quantization of large language models

by equivalent and effective shifting and scaling. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1648–1665, 2023. [2](#)

- [51] 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*, pages 38087–38099. PMLR, 2023. [2](#), [8](#)
- [52] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025. [3](#)
- [53] Zhewei Yao, Xiaoxia Wu, Cheng Li, Stephen Youn, and Yuxiong He. Exploring post-training quantization in llms from comprehensive study to low rank compensation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 19377–19385, 2024. [1](#), [2](#), [6](#)