

OWQ: Outlier-Aware Weight Quantization for Efficient Fine-Tuning and Inference of Large Language Models

Changhun Lee^{1*}, Jungyu Jin^{2*}, Taesu Kim³, Hyungjun Kim³, Eunhyeok Park²

¹Department of Convergence IT Engineering, POSTECH

²Graduate School of Artificial Intelligence, POSTECH

³SqueezeBits Inc.

{changhun.lee, jgjin0317}@postech.ac.kr, {taesu.kim, hyungjun.kim}@squeezebits.com, eh.park@postech.ac.kr

Abstract

Large language models (LLMs) with hundreds of billions of parameters require powerful server-grade GPUs for inference, limiting their practical deployment. To address this challenge, we introduce the outlier-aware weight quantization (OWQ) method, which aims to minimize LLM's footprint through low-precision representation. OWQ prioritizes a small subset of structured weights sensitive to quantization, storing them in high-precision, while applying highly tuned quantization to the remaining dense weights. This sensitivity-aware mixed-precision scheme reduces the quantization error notably, and extensive experiments demonstrate that 3.1-bit models using OWQ perform comparably to 4-bit models optimized by OPTQ. Furthermore, OWQ incorporates a parameter-efficient fine-tuning for task-specific adaptation, called weak column tuning (WCT), enabling accurate task-specific LLM adaptation with minimal memory overhead in the optimized format. OWQ represents a notable advancement in the flexibility, efficiency, and practicality of LLM optimization literature. The source code is available at <https://github.com/xvyaward/owq>.

Introduction

Large language models (LLMs) (Brown et al. 2020; Radford et al. 2019; Scao et al. 2022; Touvron et al. 2023; Zhang et al. 2022) demonstrate impressive generation performance on a wide range of complex language tasks, triggering explosive growth in LLM-based applications. However, the extensive memory and computational demands are major obstacles to the widespread use of LLMs, not only for training but for inference as well. For instance, using fp16, the GPT3-175B model necessitates approximately 330 GB of space merely to store model parameters, which eventually costs hundreds of thousands of dollars to build the system with multiple server-grade GPUs. For the widespread adoption of LLMs, it is crucial to minimize such serving costs.

Recently, weight quantization has emerged as an attractive optimization method for LLMs (Frantar et al. 2023). By storing parameters into low-precision representation, storage space can be considerably saved, and this also introduces performance benefits by addressing memory bottlenecks and

reduced communication costs (Park et al. 2023). Advanced studies (Frantar et al. 2023; Park et al. 2023) have shown that matrix multiplication with 3-bit weight and fp16 activation exhibits remarkable performance improvements on a single GPU compared to the case with fp16 weight and activation on multiple GPUs. Weight quantization of LLMs could resolve the memory and performance issues of LLMs jointly.

Previously, OPTQ (Frantar et al. 2023), also known as GPTQ (Frantar et al. 2022), introduced a layer-wise post-training quantization (PTQ) method based on the optimal brain compression (OBC) algorithm (Frantar and Alistarh 2022). Notably, the compressed 3-bit OPT-175B model outperforms the fp16 OPT-30B model, even though they have similar memory footprints. The compressed model now occupies around 63 GB, allowing deployment on a single A100 GPU. However, there is room for improvement. OPT-175B still experiences some degradation with 3-bit quantization, and this effect is more pronounced in smaller models. Given that various model sizes are optimal for different scenarios, maintaining their accuracy remains highly important.

On the other hand, weight quantization broadens its application to task-specific fine-tuning. For example, QLoRA (Dettmers et al. 2023) enables the fine-tuning of the quantized LLMs for target tasks by incorporating low-rank high-precision tensors into the quantized dense matrix. During this process, only the added tensors are updated, while the quantized path remains unchanged. This method is gaining attention as it allows for efficient fine-tuning with reduced memory consumption while mitigating the drawback, quantization errors, during fine-tuning. However, the quality of the low-precision dense matrix is often overlooked, even though its suboptimal quality could diminish the benefits of fine-tuning. Even for the adaptation, high-quality weight quantization is essential.

In this paper, we introduce a new weight quantization technique, Outlier-aware Weight Quantization (OWQ). OWQ is specifically designed considering the unique characteristics of LLMs, the presence of activation outliers (Dettmers et al. 2022; Wei et al. 2022b; Xiao et al. 2022). Our analysis reveals that these outliers play a key role in the quality degradation of weight quantization. Based on this finding, we design OWQ that applies a mixed-precision quantization considering the sensitivity of each weight column. Extensive analysis indicates that the 3.1-bit OWQ

*These authors contributed equally.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

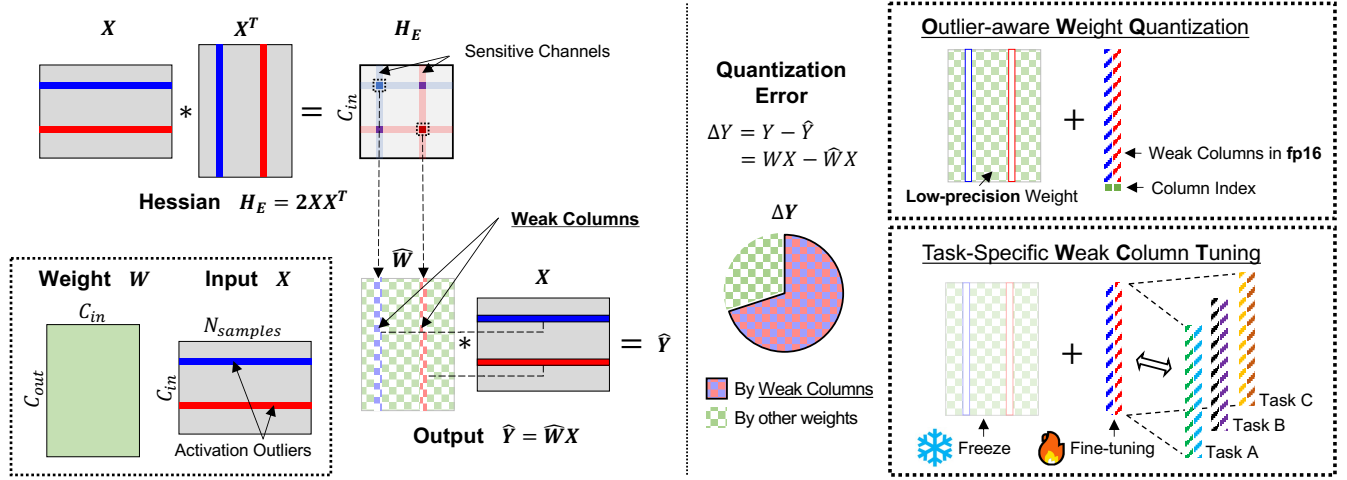


Figure 1: The overview of the proposed weak columns concept, Outlier-aware Weight Quantization (OWQ) scheme, and Weak Column Tuning (WCT) scheme for efficient task-specific fine-tuning.

model has comparable quality to the 4-bit OPTQ model. Moreover, we introduce an effective fine-tuning method based on OWQ, called Weak Column Tuning (WCT); it shares the benefit of OWQ during fine-tuning and inference, so the network can be updated to the target task with minimal memory overhead and accelerated during inference. Furthermore, the WCT-based fine-tuning model adapts with fewer trainable parameters than existing methods and outperforms them because of the superior representation quality of the weights quantized by OWQ. To our knowledge, this is the first study to consider the existence of activation outliers in extremely low-precision weight quantization and closely integrate it with fine-tuning.

Background and Related Works

Quantization and LLMs

Quantization is a widely used optimization technique aimed at exploiting the benefit of low-precision while maintaining the quality of the network. While its primary benefit is size reduction, quantization can also substantially accelerate performance through support for low-precision operations. However, a trade-off exists: quantization can lead to quality degradation, which numerous studies have aimed to address. Early studies focused on quantization-aware training (QAT) (Esser et al. 2019; Zhou et al. 2016), which tried to restore quality through additional training; however, as the understanding of quantization grew and various techniques emerged, post-training quantization (PTQ) (Li et al. 2021; Nagel et al. 2019, 2020; Wei et al. 2022a) have been actively studied, enabling quality preservation without training.

Due to the LLM’s necessity of significant storage space and computational resources, it is crucial to apply optimization via quantization. In general, QAT has been favored for extremely low-bit precision to minimize quantization error. However, it is less favorable for LLM quantization because of the high cost of the training environment. Instead, PTQ has emerged as an important topic for LLM quantization.

This field has two distinct approaches: one aims to quantize both activations and weights to int8 (Dettmers et al. 2022; Xiao et al. 2022), considering both capacity reduction and acceleration. In contrast, the second approach focuses solely on weight quantization to sub-4-bit precision (Frantar et al. 2023; Park et al. 2023). In this paper, we align our work with the latter approach. While concentrating on weight quantization, we devise a novel quantization scheme, drawing significant inspiration from int8-related research on activation quantization.

Int8 Quantization for Activation and Weight

Int8 multiplication can provide up to 2x performance improvements and more than 5x energy consumption reduction compared to fp16 baselines (Horowitz 2014). Numerous studies (Dettmers et al. 2022; Xiao et al. 2022) aim to quantize both activation and weight to int8 for matrix multiplication operations in LLMs. However, those studies identify a unique challenge of LLMs for activation quantization. LLMs exhibit a few outliers in intermediate activations, with values significantly larger than others, and these outliers are concentrated in specific feature dimensions. Preserving the values of these outliers is known to be crucial for maintaining accuracy after activation quantization.

In this study, while only weight quantization is applied, we figure out that the presence of activation outliers still impacts the sensitivity of weight quantization. We also demonstrate that considering activation outliers is essential for accurate weight quantization.

OPTQ: Weight Quantization for LLMs

OPTQ (Frantar et al. 2023) is the state-of-the-art research in the field of weight quantization for LLMs. It is based on Optimal Brain Compression (OBC) (Frantar and Alistarh 2022), which employs element-wise quantization (pruning) and compensation, using a Hessian-based metric of layer-wise quantization errors (Eq. (1) and Eq. (2)). This approach

differs from previous studies, which applied quantization through gradient descent based on the straight-through estimator (Li et al. 2021; Wei et al. 2022a) or the rounding-to-nearest mechanism (Dettmers et al. 2022).

$$w_q = \underset{w_q}{\operatorname{argmin}} \frac{(\operatorname{quant}(w_q) - w_q)^2}{[H_F^{-1}]_{qq}}, \quad (1)$$

$$\delta_F = -\frac{w_q - \operatorname{quant}(w_q)}{[H_F^{-1}]_{qq}} \cdot (H_F^{-1})_{:,q} \quad (2)$$

OPTQ has optimized OBC to parallelize quantization for each element of the output channel dimension, enabling rapid quantization. While it showcases the potential of sub-4-bit quantization for LLMs, reducing the model size or increasing the problem’s complexity results in decreased accuracy compared to the fp16 baselines. In this paper, we suggest selectively applying high-precision to weights that are vulnerable to quantization caused by activation outliers and applying the OPTQ to the remaining weights with the modification based on quantization configuration tuning for additional error reduction. These enhancements can significantly reduce the quantization error while preserving the quantization speed of OPTQ.

Parameter-Efficient Fine-Tuning (PEFT)

Fine-tuning LLMs for specific tasks can improve performance on unseen or complex tasks, but the large number of parameters requires a hyperscale computation system, which is often impractical due to high costs. Therefore, Parameter-Efficient Fine-Tuning (PEFT) schemes have been introduced to address this issue. LoRA (Hu et al. 2022) exemplifies PEFT by freezing pre-trained weights but incorporating a small fraction of learnable parameters through low-rank decomposition. As only the added parameters are updated, LoRA can be adapted using significantly less memory. QLoRA (Dettmers et al. 2023) further reduces memory consumption by replacing dense weights with quantized weights, making the fine-tuning process more lightweight. Since fine-tuning can mitigate quantization errors, QLoRA emerges as a desirable optimization for the practical, task-specific deployment of LLMs. In this study, we introduce an OWQ-compatible PEFT scheme. The superior representation quality of OWQ yields exceptional performance after PEFT with lower resource overhead than QLoRA.

Problem Definition and Motivation

In this section, before introducing our idea, we first define the problem and explain our findings clearly. The proposed OWQ is designed to apply layer-wise uniform quantization for the weights of LLMs with minimal quality degradation. When an input feature $X \in R^{C_{in} \times N}$ is given, where C_{in} represents the number of input channels and N is the sequence length of the input, the full-precision weight matrix $W \in R^{C_{out} \times C_{in}}$ for C_{out} output features are mapped to low-precision toward minimizing the difference of output activations before and after quantization. The objective function to find the quantized weight \hat{W} that minimizes the

squared error is defined as follows:

$$\arg \min_{\hat{W}} E = \arg \min_{\hat{W}} \|WX - \hat{W}X\|_2^2 \quad \text{s.t.} \quad C(\hat{W}) < C_t, \quad (3)$$

where $C(\cdot)$ represents the compression ratio and C_t is the target compression ratio. The layer-wise quantization process is applied sequentially from the model input to the output, ensuring the quantization of all weights in the model. We keep embedding and head weights with full-precision.

Layer-wise Quantization and Hessian of Weights

In this subsection, we explain our insight concerning the relationship between weight sensitivity and activation outliers in the context of quantization: weights linked to activation outliers are particularly susceptible to quantization. This understanding serves as the core motivation of OWQ.

Initially, we restructure the squared error term in Eq. (3) to represent the sum of squared errors for each output channel within the weight matrix, resulting in the equation $\sum_{i=1}^{C_{out}} \|W_{i,:}X - \hat{W}_{i,:}X\|_2^2$. This decomposition distinctly showcases that the overall error is divided into individual errors for each output channel. With the modified equation, our focus shifts to two key aspects. Firstly, it is important to note that there is no Hessian interaction between output channels. Specifically, the individual Hessians with respect to the layer-wise quantization error, denoted as $H^{(i)} \in R^{C_{in} \times C_{in}}$, have an identical value as:

$$H^{(i)} = H = \frac{\partial^2 E_i}{\partial W_{i,:}^2} = 2XX^T. \quad (4)$$

Secondly, as observed in previous studies (Nagel et al. 2020), the individual error term can be approximated using Taylor expansion. By setting $\Delta W_{i,:} = W_{i,:} - \hat{W}_{i,:}$, the error for the i -th output channel can be expressed as follows:

$$E_i = \|W_{i,:}X - \hat{W}_{i,:}X\|_2^2 \approx \Delta W_{i,:}H\Delta W_{i,:}^T. \quad (5)$$

The equation shows that in the context of layer-wise quantization, the output error can be directly related to the Hessian and the magnitude of weight perturbation.

Keeping these observations in mind, we can derive an interesting insight by acknowledging the presence of activation outliers in LLMs. Previous studies (Dettmers et al. 2022; Xiao et al. 2022) have reported that certain feature dimensions of LLM activation contain outliers with significantly larger values than others. As shown in Figure 1 Left, these activation outliers make some elements of H have exceptionally high values. This abnormal surge in Hessian values increases the corresponding weight channels’ sensitivity to quantization. In detail, as indicated in Eq. (5), even when the same degree of weight perturbation is present, the ensuing change in output can be considerably larger due to some large elements of H . We refer to the weights susceptible to quantization as weak column, specifically those associated with the activation outliers in a specific input channel.

Therefore, if we quantize all weights to the same bit-width during the weight quantization process, the quantization error from the weak columns corresponding to the activation

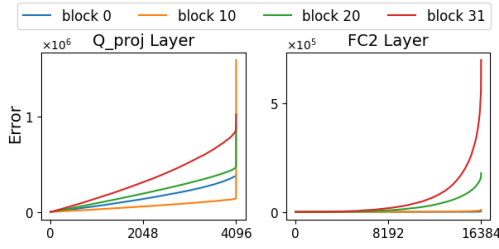


Figure 2: Cumulative error for each input channel. The x-axis channel indices are sorted in ascending order based on their influence on the output error.

outliers can lead to substantial perturbation on the output, resulting in a notable quantization error. Figure 2 supports this assertion, indicating that a large portion of the error originates from a limited number of channels, which align with the weak columns. To minimize the error from weak columns, it’s imperative to specially address these columns.

OWQ: Outlier-aware Weight Quantization

To tackle this issue, we introduce a novel concept termed Outlier-aware Weight Quantization (OWQ). OWQ encompasses two steps: initially, it identifies the weak columns and excludes them from quantization. Subsequently, it quantizes the remaining weights to extreme low-precision using meticulously tuned quantization parameters. The overview of the proposed OWQ scheme is illustrated in Figure 1 Left. In this section, we thoroughly discuss the details of the OWQ implementation.

In Eq. (4), we highlighted the relationship between the Hessian matrix and sensitivity caused by activation outliers. We also demonstrated that the final error is influenced by the quadratic terms of perturbations with the Hessian matrix in Eq. (5). Building on these insights, we define the sensitivity of j -th weight column as follows:

$$\text{sensitivity}_j = \lambda_j \|\Delta W_{:,j}\|_2^2, \quad (6)$$

where λ_j is the j -th diagonal element of the Hessian matrix. By analyzing the sensitivity of individual columns, we can effectively identify the weak columns that are vulnerable to quantization and require higher precision. When the goal is to select a specific number (k) of weak columns, the proposed metric is utilized to choose the top- k sensitive columns based on their sensitivity values.

It is worth noting that hessian-based metrics have been often used in previous studies (Dong et al. 2019, 2020). While our work has a distinctive difference in the granularity of the quantization domain and detailed expression, our observation is well aligned with the intuition of the existing studies.

Following the selection of weak columns, the remaining weights are quantized into low-precision. Any low-precision quantization scheme is applicable, but we employ OPTQ. Since OPTQ also utilizes sequential column-wise quantization, the weights excluding the weak columns can be seamlessly integrated into the OPTQ framework. Additionally, please note that the weak columns can be used to further mitigate errors that occur during the OPTQ process. OPTQ

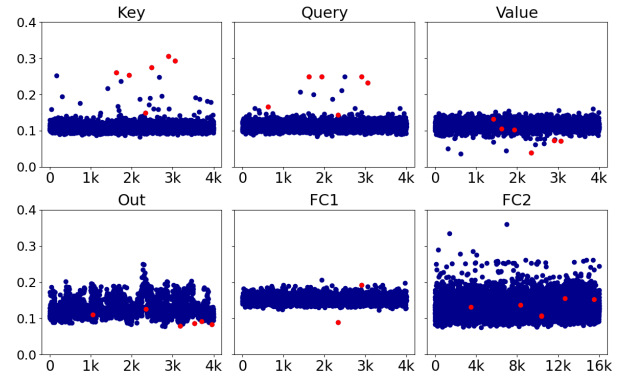


Figure 3: Min-max range of the weights in each column (blue dots) and the selected weak columns (red dots).

updates the remaining unquantized weights to compensate for the errors caused by quantizing the weight columns in the current step, as shown in Eq. (2). By rearranging the high-precision weak columns to the end of the weight before utilizing OPTQ, quantization errors from other columns during the OPTQ process can be largely compensated for by the weak columns. Since the weak columns are retained with full-precision, these compensated values can be preserved even if all other columns are quantized.

After identifying the weak columns and quantizing the remaining weights, we store the weak columns as fp16 and use an extra single integer per column, which represents the column index of the weak columns. In addition, we store a low-precision matrix whose positions of weak columns are zero-filled. Therefore, compared to OPTQ, the additional storage overhead is solely caused by the weak columns. This overhead is negligible ($\approx 0.3\%$), while the accuracy is significantly improved. In addition, we also provide the specialized acceleration for OWQ format on real GPU. A comprehensive explanation of the implementation will be provided in the “Acceleration on Real Device” section.

Additionally, as depicted in Figure 3, our approach is distinctly different from existing outlier-aware quantization studies (Park, Kim, and Yoo 2018; Park, Yoo, and Vajda 2018). While the naming is similar, the outcome is completely different. In those studies, outlier weights are excluded from quantization based solely on their magnitude to minimize the error of weights before and after the quantization. However, our method minimizes the error of output activation. As shown in the figure, the weights are selected based on their sensitivity rather than their magnitude.

Quantization Configuration Search

We mentioned that we use OPTQ for the quantization of the dense weights except for the weak columns. However, we made an important modification to OPTQ to further reduce quantization error. While OPTQ originally relies on a straightforward min-max quantization approach, we have incorporated the benefits of truncation. Many previous studies (Esser et al. 2019; Li et al. 2021; Nahshan et al. 2021; Wei et al. 2022a) pointed out that the truncation reduces the

overall quantization error significantly, improving the quality of output significantly. In OWQ, we searched the quantization configurations, such as step size and zero point using a simple 2D grid search, to narrow the quantization range. The optimal values of the parameters that minimize the difference before and after quantization are searched via rounding to the nearest with truncation, and we apply OPTQ on top of the searched values.

This modification substantially reduces the overall reconstruction error, significantly improving quality. In OWQ, the WikiText-2 perplexity is lowered from 12.14 to 11.21 for OPT-6.7B. An interesting observation is that when truncation is applied to conventional OPTQ, output quality degrades, with the WikiText-2 perplexity increasing from 12.88 to 48.26 in OPT-6.7B. We have observed that the weak columns in the key and query layers of transformer blocks have exceptionally large values. Truncating these leads to substantial error, resulting in a significant reduction in accuracy. However, with our approach, this error is avoided because weak columns are retained as full-precision. We argue that this enhancement is a significant advantage of OWQ, notably improving the representation quality of the low-precision matrix.

PEFT with Weak Column Tuning

When applied to pre-trained LLMs, OWQ results in minimal quality degradation, allowing the model to maintain its zero-shot generative performance. Nonetheless, to further boost output quality, fine-tuning might be necessary for newly introduced tasks. In this context, we present a Weak Column Tuning (WCT) scheme, illustrated in Figure 1 Right. WCT first quantizes the base model with OWQ and then fine-tunes only the weak columns that retained high-precision as a result of OWQ. In OWQ, we identify weak columns using channel-wise sensitivity and retain them in high-precision, meaning slight alterations in these sensitive columns have a substantial impact on the output. Moreover, because they are high-precision (fp16), the values can be freely modified. Empirically, we’ve found that updating these weak columns facilitates the adaptation of the quantized network to the desired task. It’s worth noting that the low-precision dense matrix remains static during this update.

Because learnable weak columns represent only a small fraction of the overall parameters, updates also result in negligible memory overhead. Furthermore, the OWQ format is preserved after fine-tuning, enabling us to leverage the benefits of acceleration with the customized kernel for OWQ. Additionally, by allocating per-task weak columns, we can serve multiple task-specific models with minimal memory overhead through shared utilization of the dense low-precision matrix, which constitutes the majority of the memory footprint.

The versatile adaptability of OWQ broadens its range of applications. Furthermore, as will be shown in the “Experiment” section, WCT surpasses leading fine-tuning methods in terms of both memory usage and output quality. This is because OWQ improves the quality of dense low-precision representation, and sensitive columns are pre-picked and used as the update target. These factors enhance the qual-

ity of the fine-tuned network significantly. Please refer to the “Comparison of PTQ Methods used in WCT” section for details.

Experiments

Experimental Setup

To validate the outstanding performance of our proposed method, we present quantization results for large-scale LLMs such as OPT (Zhang et al. 2022) and LLaMA (Touvron et al. 2023) families. Our primary baseline is OPTQ, so we apply identical experimental settings of it. For instance, our calibration dataset consists of 128 random 2048 token segments from the C4 dataset (Raffel et al. 2020). Experiments were conducted on a single NVIDIA A100 GPU with 80 GB of main memory or RTX3090 GPU with 24 GB memory. Like OPTQ, our method quantizes the target model without re-training. To measure the zero-shot or few-shot performance, we utilize an open-source evaluation repository, EleutherAI/lm-evaluation-harness (Gao et al. 2023). Please note that we report the numbers with an error margin based on 3 experiments with different seeds.

Compared to the baselines of 3-bit and 4-bit OPTQ results, we present two variants, with an extra 0.01 bit and 0.1 bit overhead, respectively. For the 4-bit baseline, we found that 0.01 bit overhead is sufficient. The additional storage area of extra bit is evenly distributed across the linear layers within the transformer block. For instance, in the OPT model which has six linear layers (key, query, value, out, fc1, and fc2), the weak columns of the key layers contribute an average of 0.00167 bit in the 3.01-bit configuration (0.15% columns of the key weight matrix). This extra bit covers the all overhead of the mixed-precision representation. If we quantize OPT-175B model with an average of 3.01 bits, it will require approximately 260 MB of additional storage compared to the 3-bit OPT-175B OPTQ model, which utilizes around 63.1 GB of storage space. All experiments were conducted using the PyTorch 2.0 (Paszke et al. 2019) framework with HuggingFace integration (Wolf et al. 2019).

While not discussed in OPTQ papers, the “act-order” (AO) option was recently added to their official GitHub. This option quantizes columns in order according to the activation magnitude and is distinct from the earlier OPTQ approach that quantized columns sequentially. Similarly, the “true-sequential” (TS) option was introduced, which applies quantization sequentially, taking into account the quantization error from previous layers within a block. Both methods boost the accuracy of OPTQ, and for a comprehensive comparison, we provide OPTQ results with “TS + AO” configuration by default. We don’t use these options for OWQ unless otherwise noted.

Results of Perplexity Measure

The accuracy of the proposed model is assessed through the evaluation on multiple language tasks, including WikiText-2 (Merity et al. 2016), Penn Treebank (PTB) (Marcus et al. 1994), and C4. Perplexity-based tasks are particularly sensitive to model quantization (Frantar et al. 2023), with perplexity numbers serving as indicators of the generative per-

OPT	Bits	125M	350M	1.3B	2.7B	6.7B	13B	30B	66B
full	16	27.65	22.00	14.63	12.47	10.86	10.13	9.56	9.34
RTN	4	37.28	25.94	48.17	16.92	12.10	11.32	10.98	110
OPTQ	4	32.05	23.87	15.47	12.83	11.14	10.29	9.57	9.34
OWQ	4.01	29.47* _{±.17}	23.19* _{±.16}	15.01 _{±.06}	12.39 _{±.02}	10.87 _{±.02}	10.26 _{±.02}	9.50 _{±.01}	9.25 _{±.04}
OPTQ	3	53.43	32.28	20.90	16.55	12.88	11.58	10.29	9.90
OWQ	3.01	35.26* _{±.61}	26.59* _{±.51}	16.40 _{±.15}	13.21 _{±.09}	11.21 _{±.05}	11.48 _{±.03}	9.61 _{±.02}	9.28 _{±.03}
OWQ+TS+AO	3.01	35.19* _{±.68}	25.86* _{±.25}	15.96 _{±.18}	13.28 _{±.06}	11.20 _{±.05}	11.14 _{±.02}	9.66 _{±.06}	9.31 _{±.01}
OWQ	3.1	33.41 _{±.25}	26.00 _{±.14}	15.39 _{±.06}	12.98 _{±.10}	11.14 _{±.03}	10.38 _{±.01}	9.57 _{±.02}	9.30 _{±.02}

Table 1: OPT WikiText-2 perplexity (\downarrow). For the results with *, we used an extra 0.05 bits instead of 0.01 bits; there are few or no weak columns in the budget of 0.01 bits due to the small model dimension. TS: True-sequential and AO: Act-order options.

LLaMA	Bits	7B	13B	30B	65B
full	16	5.68	5.09	4.10	3.53
RTN	4	6.29	5.53	4.54	3.92
OPTQ	4	6.10	5.36	4.45	4.10
OWQ	4.01	5.94 _{±.02}	5.25 _{±.00}	4.25 _{±.00}	3.74 _{±.02}
OPTQ	3	8.13	6.67	5.67	5.41
OWQ	3.01	6.66 _{±.04}	5.66 _{±.01}	4.75 _{±.02}	4.25 _{±.01}
OWQ	3.1	6.41 _{±.01}	5.56 _{±.01}	4.63 _{±.01}	4.09 _{±.01}

Table 2: LLaMA WikiText-2 perplexity (lower is better).

formance of the quantized model. The results for WikiText-2 can be found in Table 1 and Table 2.

The results clearly demonstrate that OWQ consistently delivers substantial quality improvements across the LLM families, irrespective of the model size. The 3.01-bit OWQ model effectively mitigates the quality degradation observed in the 3-bit OPTQ model, while the 3.1-bit model achieves comparable performance to the 4-bit OPTQ model. Furthermore, OWQ 4.01-bit yields noteworthy improvements, highlighting the significance of treating weak columns. These results underscore the importance and effectiveness of our approach in preserving model quality after quantization.

An interesting finding is the significant improvement in model quality for OPT models with less than 13 billion parameters when applying OWQ. Although previous studies have highlighted the presence of activation outliers in models with more than 6.7 billion parameters (Dettmers et al. 2022), even smaller models with moderately large channels can still benefit from mixed precision quantization. This suggests that the concept of weak columns remains valid and effective in enhancing the quality of LLMs, regardless of the model size.

In addition, please note that OWQ consistently outperforms OPTQ with TS + AO, but TS + AO doesn't give performance benefits to OWQ (Table 1). Although no theoretical interpretation for those options was proposed, our study suggests that the benefits of "act-order" arise from sensitivity-aware quantization. This means that applying sequential quantization beginning with sensitive columns improves performance for OPTQ. However, act-order alone cannot sufficiently mitigate the quality degradation caused

by weak columns within a low-precision domain. Interestingly, the benefit of TS is also mitigated as quantization error is already greatly reduced by using OWQ.

Results of Various Few-shot Tasks

We conducted additional experiments on diverse few-shot language tasks. Referring to the "Open LLM Leaderboard" (Beeching et al. 2023) from the Huggingface H4 team as a benchmark, we relied on the average scores from ARC-challenge (25-shot) (Clark et al. 2018), Hellaswag (10-shot) (Zellers et al. 2019), and MMLU (5-shot) (Hendrycks et al. 2021). In particular, MMLU is a collection of 57 tasks. These benchmarks were chosen because they test a range of reasoning and general knowledge across diverse fields, in few-shot contexts. The average scores of Table 3 and Table 4 convince us that the proposed OWQ is consistently superior to OPTQ in various model sizes, from 125m to 66B. In particular, we can see the considerable gap between OPTQ 3-bit and OWQ 3.01-bit for the LLaMA families in Table 4 despite using the TS and AO options for OPTQ. Our method's strength lies in its universality, consistently boosting the performance of generative models with minimal storage overhead.

Acceleration on Real Device

To demonstrate the advantages of low-precision acceleration, we developed a customized CUDA kernel for OWQ and assessed its latency overhead on an A100 GPU. First, we decompressed the low-precision matrix into the fp16 format and performed dense GeMV multiplication. At this stage, the overhead is identical to that of OPTQ's customized kernel. Additionally, we select the activation input channels for these weak columns on-the-fly and use another GeMV kernel specifically for them. This approach effectively avoids the issues of irregular memory access. Table 5 displays the kernel overhead of OWQ 3.01-bit for various model sizes. The mixed-process computation adds up to only 3.21% latency compared to the 3-bit acceleration of the OPTQ kernel on the LLaMA 7B model, and the overhead is generally amortized for larger models.

OPT	Bits	125M	350M	1.3B	2.7B	6.7B	13B	30B	66B
full	16	26.78	28.77	36.34	40.27	44.22	45.34	47.97	49.86
OWQ	4.01	26.67* $\pm.29$	28.85* $\pm.11$	36.48 $\pm.19$	39.75 $\pm.05$	43.94 $\pm.03$	45.42 $\pm.14$	47.83 $\pm.22$	49.57 $\pm.07$
OPTQ	4	26.43	28.31	36.41	39.66	43.69	44.97	47.58	49.47
OWQ	3.1	25.99 $\pm.28$	27.77 $\pm.23$	35.86 $\pm.13$	38.93 $\pm.19$	42.80 $\pm.07$	44.50 $\pm.22$	47.20 $\pm.28$	49.02 $\pm.25$
OWQ	3.01	26.18* $\pm.40$	27.91* $\pm.17$	35.51 $\pm.24$	38.79 $\pm.81$	42.51 $\pm.26$	44.62 $\pm.13$	47.17 $\pm.20$	48.84 $\pm.03$
OPTQ	3	26.00	27.78	34.01	37.16	41.63	43.70	46.65	48.01

Table 3: The average value of three few-shot scores (%): ARC-challenge, Hellaswag, and MMLU for OPT families (higher is better). For the results with *, we used an extra 0.05 bits instead of 0.01 bits due to the small model dimension.

LLaMA	Bits	7B	13B	30B	65B
full	16	54.81	61.56	68.21	71.17
OWQ	4.01	54.01 $\pm.04$	60.54 $\pm.52$	68.12 $\pm.15$	70.30 $\pm.06$
OPTQ	4	52.91	59.90	66.96	69.92
OWQ	3.1	51.93 $\pm.48$	58.84 $\pm.36$	66.18 $\pm.11$	69.26 $\pm.29$
OWQ	3.01	50.21 $\pm.45$	57.69 $\pm.21$	66.21 $\pm.25$	68.82 $\pm.29$
OPTQ	3	47.06	52.87	61.15	65.75

Table 4: The average value of three scores (%): ARC-challenge, Hellaswag, and MMLU for LLaMA families.

OPT	6.7B	13B	66B	LLaMA	7B	13B	65B
k/q/v/o	3.07	3.11	2.84	k/q/v/o	3.21	3.47	2.34
fc1	2.31	1.53	1.65	up/gate	2.34	2.01	1.65
fc2	2.15	2.04	2.23	down	2.03	2.21	2.03

Table 5: Kernel overhead of OWQ 3.01-bit (vs OPTQ) (%).

Quantization Speed

For LLM quantization, the quantization algorithm’s speed is crucial. While OWQ adds operations for weak column selection and hyperparameter tuning compared to OPTQ, sharing the Hessian with OPTQ minimizes OWQ’s overhead. Furthermore, applying the true-sequential option to OPTQ adds extra runtime for OPTQ, further reducing the gap with OWQ quantization time. On A100 GPU, OWQ can quantize a 66B model in under 3 hours, presenting its practicality.

Results of WCT-based Fine-tuning

To validate the superior performance of Weak Column Tuning (WCT) for task-specific adaptation, we fine-tuned the LLaMA 7B model with 4-bit quantization using WCT and then compared the results. To obtain the results of QLoRA (Dettmers et al. 2023), we used the official checkpoint. For a fair comparison, we used the same subset of the OpenAssistant (Köpf et al. 2023) dataset that QLoRA employed. We evaluated performance by inputting the results generated by both models for 80 questions from the Vicuna Benchmark (Chiang et al. 2023) into GPT-4 (OpenAI 2023) to determine which was better. As reported in a previous work (Dettmers et al. 2023), we found a bias in which GPT-4 favors the system that appears first in a given pair of systems, giving it a

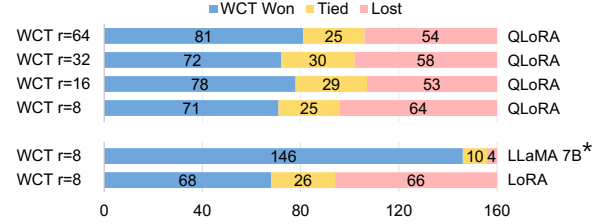


Figure 4: GPT-4 based analysis of WCT-based fine-tuning. * denotes the base model (without fine-tuning).

higher score. To eliminate this bias, we evaluate both ordering cases and report results for a total number of 160 evaluations. LoRA (Hu et al. 2022) and QLoRA both use the 64 rank of adapter modules. We used nucleus sampling with $p = 0.9$ and temperature 0.7 for all generations.

In the WCT experiments, $r = k$ means the configuration with r weak columns for each layer. As depicted in Figure 4, WCT with 64 weak columns ($r = 64$) surpasses the QLoRA. In other words, GPT-4 evaluated the tuning results using WCT as better more often (81 vs 54 for $r = 64$ case). Remarkably, WCT with just 8 weak columns (takes only 6.8% of learnable parameters vs. QLoRA) outperforms QLoRA and yields results comparable to full-precision LoRA. As the quality of the compressed weight is on par with the full-precision model, OWQ combined with WCT delivers performance matching that of full-precision LoRA only with 24.4% of overall memory usage during inference. WCT only updates weak columns, which are highly sensitive to update. This feature makes WCT compensate accuracy with a smaller rank than the conventional LoRA. With OWQ + WCT, we can enjoy the benefits of quantization not only in inference but also in task-specific adaptation.

Comparison of PTQ Methods used in WCT

We compare several quantization methods used for quantizing fixed dense weights in WCT and verify that the sophisticated quantization method is important for performance after fine-tuning. After quantizing the LLaMA-7B model with different quantization methods, the models refined with WCT are asked to generate answers to questions from the Vicuna Benchmark, and the quality of the answers is compared using GPT-4-based evaluation. For all quantization

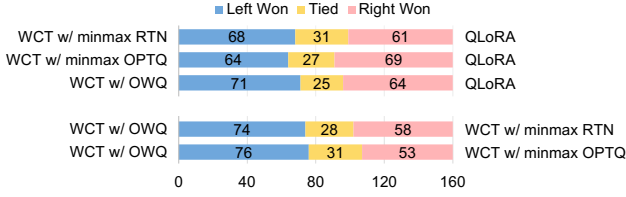


Figure 5: Comparison of fine-tuned performance of several post-training quantization methods used for WCT.

Method / Bits / Group	OPT-6.7B			
	Wiki2	PTB	C4	lamb. \uparrow
OPTQ / 3	12.76	19.35	14.55	60.15
OPTQ / 3 / g1024	12.05	18.01	13.89	64.44
OPTQ / 3 / g128	11.55	17.26	13.43	64.96
OWQ / 3.01	11.22	16.32	13.23	69.94
OWQ / 3.01 / g1024	11.18	16.32	13.19	68.91
OWQ / 3.01 / g128	11.16	16.23	13.10	68.61

Method / Bits / Group	LLaMA 7B			
	Wiki2	PTB	C4	winog. \uparrow
OPTQ / 3	8.08	14.13	10.26	64.33
OPTQ / 3 / g1024	7.15	12.75	9.12	66.26
OPTQ / 3 / g128	6.56	12.48	8.36	67.60
OWQ / 3.01	6.65	12.47	8.62	67.05
OWQ / 3.01 / g1024	6.61	12.05	8.49	67.60
OWQ / 3.01 / g128	6.40	11.72	8.18	69.57

Table 6: The results of OWQ and OPTQ with group-wise quantization. lamb. = lambada, winog. = winogrande.

methods, we used $r = 8$ as the number of weak columns and linear asymmetric quantization with the per-channel granularity. Although the performance loss due to quantization can be compensated by fine-tuning, Figure 5 shows that the quality of the quantized model before fine-tuning affects the performance after fine-tuning.

Comparison with Group-wise Quantization

Applying uniform quantization at fine-grained granularity significantly reduces quantization error while introducing some storage overhead for quantization hyperparameters. OPTQ utilizes this approach by dividing row vectors into groups (e.g., group size of 128 or 1024) and applying uniform quantization independently with different configurations. This expansion can be applied orthogonally to OWQ, so we can combine it with OWQ to assess any improvements. Results in Table 6 show that the improvement from fine-grained quantization is negligible, as OWQ already substantially enhances the 3-bit model’s quality. Moreover, compared to grouped OPTQ with 128 group size, 3.01-bit OWQ’s storage overhead is only about 10% of grouped OPTQ overhead while achieving comparable or better perplexity and zero-shot accuracy. Thus, OWQ is a superior solution to the grouping technique.

Equation	$\lambda_j \Delta W_{:,j} _2^2$	λ_j	$ \Delta W_{:,j} _2^2$	$\Sigma W_{:,j} $
OPT-6.7B	11.23	11.25	13.08	14.77
LLaMA 7B	6.64	6.76	11.57	10.12

Table 7: LLaMA WikiText-2 perplexity (lower is better).

Eff. bit	3.005	3.01	3.05	3.10	3.20	3.30
Wiki-2 (\downarrow)	11.38	11.22	11.19	11.17	11.15	11.13
Lat. (%)	2.89	2.91	4.19	6.21	6.86	8.69
Mem. (%)	0.21	0.41	2.04	3.95	7.60	10.95

Table 8: Effective bit-width sweep using OPT-6.7B model. Lat. = Latency Overhead, Mem. = Memory Overhead.

Weak Column Selection Metrics

In this paper, we propose to use both the Hessian matrix and weight perturbations for weak column selection (Eq. (6)), to minimize layer output error. Table 7 additionally presents perplexity results for various weak column selection metrics. It is obvious that the results for the 1st ($\lambda_j ||\Delta W_{:,j}||_2^2$, our proposed selection metric) and 2nd (λ_j , Hessian only) columns are considerably better than those for the 3rd ($||\Delta W_{:,j}||_2^2$, weight error only). This indicates that (1) minimizing layer-wise error is a valid objective for final accuracy, and (2) it is necessary to account for all factors contributing to output activation error, rather than simply focusing on minimizing weight error.

Varying Ratios of Weak Columns

We report the trade-offs in perplexity, latency, and memory usage for varying ratios in Table 8. These overhead results are compared to the OPTQ 3-bit baseline. With just 3.005 bits, the model already surpasses the OPTQ 3-bit performance with negligible overhead. However, as the bit width increases, the performance gain saturates while overhead keeps increasing. As activation outliers are limited to only a few dimensions, even a small number of weak columns can lead to noticeable performance improvements. In the main tables, we focused on two specific ratios, 3.01 and 3.1-bit, as they effectively illustrate the trend.

Conclusion

The presence of activation outliers has been identified as a significant challenge in LLM activation quantization. We found that even in weight quantization, activation outliers can increase the sensitivity of certain weight columns, leading to significant quality degradation in a low-precision domain. To overcome this, we introduced a novel quantization scheme, OWQ. Compared to existing 3-bit quantization methods, OWQ improves quality notably with only negligible storage and computation overhead, sustaining the benefit of low-precision compression. In addition, we introduce the WCT scheme, which enables task-specific adaptation with minimal memory overhead and shows outstanding performance. We believe that our insights will promote the widespread adoption of LLMs.

Acknowledgments

This work was supported by IITP grant (MSIT, No.2021-0-00310) and NRF grant (MSIT, RS-2023-00213611, RS-2023-00228970) funded by the Korea government.

References

- Beeching, E.; Fourrier, C.; Habib, N.; Han, S.; Lambert, N.; Rajani, N.; Sanseviero, O.; Tunstall, L.; and Wolf, T. 2023. Open LLM Leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *NeurIPS*, 33: 1877–1901.
- Chiang, W.-L.; Li, Z.; Lin, Z.; Sheng, Y.; Wu, Z.; Zhang, H.; Zheng, L.; Zhuang, S.; Zhuang, Y.; Gonzalez, J. E.; Stoica, I.; and Xing, E. P. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with 90%* ChatGPT Quality.
- Clark, P.; Cowhey, I.; Etzioni, O.; Khot, T.; Sabharwal, A.; Schoenick, C.; and Tafford, O. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv:1803.05457*.
- Dettmers, T.; Lewis, M.; Belkada, Y.; and Zettlemoyer, L. 2022. Llm.int8(): 8-bit matrix multiplication for transformers at scale. *arXiv:2208.07339*.
- Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. *arXiv:2305.14314*.
- Dong, Z.; Yao, Z.; Arfeen, D.; Gholami, A.; Mahoney, M. W.; and Keutzer, K. 2020. Hawq-v2: Hessian aware trace-weighted quantization of neural networks. *NeurIPS*, 33: 18518–18529.
- Dong, Z.; Yao, Z.; Gholami, A.; Mahoney, M. W.; and Keutzer, K. 2019. Hawq: Hessian aware quantization of neural networks with mixed-precision. In *CVPR*, 293–302.
- Esser, S. K.; McKinstry, J. L.; Bablani, D.; Appuswamy, R.; and Modha, D. S. 2019. Learned step size quantization. *arXiv:1902.08153*.
- Frantar, E.; and Alistarh, D. 2022. Optimal Brain Compression: A Framework for Accurate Post-Training Quantization and Pruning. In *NeurIPS*.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2022. GPTQ: Accurate Post-Training Quantization for Generative Pre-trained Transformers. *arXiv:2210.17323*.
- Frantar, E.; Ashkboos, S.; Hoefler, T.; and Alistarh, D. 2023. OPTQ: Accurate Quantization for Generative Pre-trained Transformers. In *ICLR*.
- Gao, L.; Tow, J.; Abbasi, B.; Biderman, S.; Black, S.; DiPofi, A.; Foster, C.; Golding, L.; Hsu, J.; Le Noac’h, A.; Li, H.; McDonell, K.; Muennighoff, N.; Ociepa, C.; Phang, J.; Reynolds, L.; Schoelkopf, H.; Skowron, A.; Sutawika, L.; Tang, E.; Thite, A.; Wang, B.; Wang, K.; and Zou, A. 2023. A framework for few-shot language model evaluation.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. *arXiv:2009.03300*.
- Horowitz, M. 2014. 1.1 computing’s energy problem (and what we can do about it). In *ISSCC*, 10–14. IEEE.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*. OpenReview.net.
- Köpf, A.; Kilcher, Y.; von Rütte, D.; Anagnostidis, S.; Tam, Z.-R.; Stevens, K.; Barhoum, A.; Duc, N. M.; Stanley, O.; Nagyfi, R.; et al. 2023. OpenAssistant Conversations—Democratizing Large Language Model Alignment. *arXiv preprint arXiv:2304.07327*.
- Li, Y.; Gong, R.; Tan, X.; Yang, Y.; Hu, P.; Zhang, Q.; Yu, F.; Wang, W.; and Gu, S. 2021. BRECQ: Pushing the Limit of Post-Training Quantization by Block Reconstruction. In *ICLR*.
- Marcus, M.; Kim, G.; Marcinkiewicz, M. A.; MacIntyre, R.; Bies, A.; Ferguson, M.; Katz, K.; and Schasberger, B. 1994. The Penn treebank: Annotating predicate argument structure. In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Merity, S.; Xiong, C.; Bradbury, J.; and Socher, R. 2016. Pointer sentinel mixture models. *arXiv:1609.07843*.
- Nagel, M.; Amjad, R. A.; Van Baalen, M.; Louizos, C.; and Blankevoort, T. 2020. Up or down? adaptive rounding for post-training quantization. In *ICML*, 7197–7206. PMLR.
- Nagel, M.; Baalen, M. v.; Blankevoort, T.; and Welling, M. 2019. Data-free quantization through weight equalization and bias correction. In *ICCV*, 1325–1334.
- Nahshan, Y.; Chmiel, B.; Baskin, C.; Zheltonozhskii, E.; Banner, R.; Bronstein, A. M.; and Mendelson, A. 2021. Loss aware post-training quantization. *Machine Learning*, 110(11-12): 3245–3262.
- OpenAI. 2023. GPT-4 Technical Report. *arXiv:2303.08774*.
- Park, E.; Kim, D.; and Yoo, S. 2018. Energy-efficient neural network accelerator based on outlier-aware low-precision computation. In *ISCA*, 688–698. IEEE.
- Park, E.; Yoo, S.; and Vajda, P. 2018. Value-aware quantization for training and inference of neural networks. In *ECCV*, 580–595.
- Park, G.; Park, B.; Kim, M.; Lee, S.; Kim, J.; Kwon, B.; Kwon, S. J.; Kim, B.; Lee, Y.; and Lee, D. 2023. LUT-GEMM: Quantized Matrix Multiplication based on LUTs for Efficient Inference in Large-Scale Generative Language Models. *arXiv:2206.09557*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1): 5485–5551.

- Scao, T. L.; Fan, A.; Akiki, C.; Pavlick, E.; Ilić, S.; Hesslow, D.; Castagné, R.; Luccioni, A. S.; Yvon, F.; Gallé, M.; et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv:2211.05100*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv:2302.13971*.
- Wei, X.; Gong, R.; Li, Y.; Liu, X.; and Yu, F. 2022a. QDrop: Randomly Dropping Quantization for Extremely Low-bit Post-Training Quantization. In *ICLR*.
- Wei, X.; Zhang, Y.; Zhang, X.; Gong, R.; Zhang, S.; Zhang, Q.; Yu, F.; and Liu, X. 2022b. Outlier Suppression: Pushing the Limit of Low-bit Transformer Language Models. In *Advances in Neural Information Processing Systems*.
- Wolf, T.; Debut, L.; Sanh, V.; Chaumond, J.; Delangue, C.; Moi, A.; Cistac, P.; Rault, T.; Louf, R.; Funtowicz, M.; et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv:1910.03771*.
- Xiao, G.; Lin, J.; Seznec, M.; Demouth, J.; and Han, S. 2022. Smoothquant: Accurate and efficient post-training quantization for large language models. *arXiv:2211.10438*.
- Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. HellaSwag: Can a Machine Really Finish Your Sentence? *arXiv:1905.07830*.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; et al. 2022. Opt: Open pre-trained transformer language models. *arXiv:2205.01068*.
- Zhou, S.; Wu, Y.; Ni, Z.; Zhou, X.; Wen, H.; and Zou, Y. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv:1606.06160*.