
GACT: Activation Compressed Training for Generic Network Architectures

Xiaoxuan Liu¹ Lianmin Zheng¹ Dequan Wang¹ Yukuo Cen² Weize Chen² Xu Han² Jianfei Chen²
Zhiyuan Liu² Jie Tang² Joseph E. Gonzalez¹ Michael W. Mahoney^{1,3,4} Alvin Cheung¹

Abstract

Training large neural network (NN) models requires extensive memory resources, and Activation Compression Training (ACT) is a promising approach to reduce training memory footprint. This paper presents GACT, an ACT framework to support a broad range of machine learning tasks for generic NN architectures with limited domain knowledge. By analyzing a linearized version of ACT’s approximate gradient, we prove the convergence of GACT without prior knowledge on operator type or model architecture. To make training stable, we propose an algorithm that decides the compression ratio for each tensor by estimating its impact on the gradient at run time. We implement GACT as a PyTorch library that readily applies to any NN architecture. GACT reduces the activation memory for convolutional NNs, transformers, and graph NNs by up to 8.1 \times , enabling training with a 4.2 \times to 24.7 \times larger batch size, with negligible accuracy loss.

1. Introduction

In recent years, we have witnessed the trend of using larger and larger neural network (NN) models to deliver improved accuracy and generalization in various machine learning tasks (Devlin et al., 2018; Fedus et al., 2021). However, training these models requires a considerable amount of on-device GPU memory. Unfortunately, the increase of GPU memory capacity has been relatively slow, leading to a fundamental barrier to the development of large NN models.

Activation Compression Training (ACT) is a promising approach to reduce the memory footprint of models during training. As all layers’ activations need to be kept in the

memory for computing the gradients during training, ACT reduces memory consumption by compressing these saved activations. Prior work (Chakrabarti & Moseley, 2019; Fu et al., 2020; Chen et al., 2021; Evans & Aamodt, 2021) has shown the effectiveness of ACT by reducing activation footprint by up to 12 \times with 2-bit activations.

Although ACT has already demonstrated impressive compression capabilities, previous work on ACT is restricted to specific NN architectures. For example, ActNN (Chen et al., 2021) is a quantization framework for convolutional NNs only; Mesa (Pan et al., 2021) proposes a per head/layer quantization method for vision transformers; and AC-GC (Evans & Aamodt, 2021) derives convergence error bound for different types of operators separately.

Developing a generic ACT framework is challenging. Theoretically, convergence guarantees must be made without assumptions on the network architecture. Algorithmically, the framework should find effective compression strategies for all kinds of networks automatically. From the system perspective, the framework should support arbitrary NN operations, including user-defined ones.

In this work, we propose GACT, a general framework for ACT that is agnostic to the NN architecture. Neither specialized mathematical derivations nor customized implementation is needed to support different operators. To enable this, we develop a general convergence theory by analyzing the stochastic gradient (SG) introduced by ACT. We show that the SG can be well approximated by a linearized version, which is unbiased to stochastic compressors. The variance of the linearized gradient has a particularly simple structure that allows a numerical algorithm to predict the variance given a compression strategy. Then, we generate the strategy by approximately solving an integer program.

We implement our method as a library based on PyTorch that can be quickly integrated into real-world machine learning systems. The library also provides several optimization levels to explore the trade-off between memory and speed. We demonstrate the flexibility and efficiency of GACT on various tasks, including image classification, object detection, text, and graph node classification. Our evaluation shows that GACT can reduce activation memory by up to 8.1 \times , enabling training with a 24.7 \times larger batch size on the

¹UC Berkeley ²Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua-Bosch Joint Center for ML, BNRist Center, State Key Lab for Intell. Tech. & Sys., Tsinghua University ³ICSI ⁴LBNL. Correspondence to: Jianfei Chen <jianfeic@tsinghua.edu.cn>.

same GPU. In sum, our main contributions are as follows:

- We propose a general convergence theory for ACT.
- We develop an algorithm that automatically estimates the sensitivity of each compressed tensor and selects the optimal compression strategy.
- We build efficient implementation of GACT in PyTorch with an easy-to-use API that can also be combined with other memory-saving techniques seamlessly.

2. Related Work

Activation Compression Training. ACT has been applied to convolutional NNs using different compressors, such as quantizers (Chakrabarti & Moseley, 2019; Fu et al., 2020; Chen et al., 2021), JPEG (Evans et al., 2020), or scientific data compression algorithms (Jin et al., 2021; Evans & Aamodt, 2021). ACT is also applied to transformers (Pan et al., 2021) and graph NNs (Anonymous, 2022).

However, the existing theory for ACT (Chakrabarti & Moseley, 2019; Fu et al., 2020; Chen et al., 2021; Evans & Aamodt, 2021) relies on the case-by-case analysis of specific network operators, such as convolution, ReLU, and batch normalization. It also requires dedicated implementations for each operator. On the contrary, GACT focuses on the generality of activation compressed training, not a specific quantizer design, which is the main topic of previous work. Instead of assuming that the network is a stack of layers, GACT formulates the problem as a computational graph of operators. This is general enough to cover transformers (Vaswani et al., 2017), graph NNs (Kipf & Welling, 2016), second-order derivatives, and unknown future architectures.

Reduced Precision Training. Apart from ACT, reduced precision training (Micikevicius et al., 2018; Wu et al., 2018; Wang et al., 2018c; Banner et al., 2018; Chen et al., 2020a; Sun et al., 2020) performs calculations directly on low precision data, reducing the computation cost and memory footprint simultaneously. To achieve this, specialized kernels are used to calculate on low precision data. In contrast, ACT only considers storage, and it can thus use more flexible compression strategies and achieve a much better compression ratio with the same accuracy loss.

Memory-Efficient Training. Gradient checkpointing (Chen et al., 2016; Jain et al., 2019) trades computation for memory by dropping some of the activations in the forward pass from memory and recomputing them in the backward pass. Swapping (Kirisame et al., 2020; Huang et al., 2020; Wang et al., 2018b; Peng et al., 2020) offloads activation or model parameters to an external memory (e.g., CPU memory). Recent work (Beaumont et al., 2021) explores the possibility of combining the gradient checkpointing and swapping. All these methods save memory by storing fewer

tensors on the GPU. In contrast, GACT compresses the saved tensors and is complementary to these approaches. Moreover, the generality of GACT enables easy combination with these methods, which we explore in this paper.

3. Formulation

We first present the mathematical formulation of our activation compressed training (ACT) framework. As we would like to develop a general ACT algorithm, applicable to a wide range of NN architectures, we make minimal assumptions on our formulation. Throughout the paper, we define the variance of a vector x as $\text{Var}[x] = \mathbb{E}[\|x\|^2] - \|\mathbb{E}[x]\|^2$.

3.1. Activation Compressed Training

In this work, we abstract the forward propagation as two functions $\ell(x; \theta)$ and $h(x; \theta)$. Both take a datum x and the model parameter θ as the input. The loss function $\ell(x; \theta)$ outputs the loss of the network θ on datum x . The context function $h(x; \theta)$ outputs tensors to be stored in the memory for computing the gradients, which are referred as the *context*. Assume that the context consists of L tensors, where each tensor $h^{(l)}(x; \theta)$ is represented by a flattened D_l -dimensional vector. Denote $h(x; \theta) = (h^{(l)}(x; \theta))_{l=1}^L$. Our notations are somewhat unconventional in the sense that we do not explicitly define each layer’s activation. We do not even assume that there is a NN. It could be any computational graph that saves context tensors.

Given a dataset $\mathcal{X} = \{x_n\}_{n=1}^N$, define the batch loss $\mathcal{L}(\theta) := \frac{1}{N} \sum_{n=1}^N \ell(x; \theta)$. The dataset can be equivalently represented as an empirical data distribution $p_{\mathcal{X}}(x) := \frac{1}{N} \sum_{n=1}^N \delta(x - x_n)$, where δ is the Dirac delta function. The batch loss can be written as $\mathcal{L}(\theta) = \mathbb{E}_{\mathcal{X}}[\ell(x; \theta)]$, where $\mathbb{E}_{\mathcal{X}}$ denotes for taking expectation over $p_{\mathcal{X}}$.

The network is trained with stochastic gradient descent (SGD) (Bottou, 2010). Starting from an initial model θ_0 , at the t -th iteration, SGD updates the model with:

$$\theta_{t+1} \leftarrow \theta_t - \eta \nabla_{\theta} \ell(x; \theta_t), \quad (1)$$

where η is a learning rate, and the SG $\nabla_{\theta} \ell(x; \theta)$ is computed on a random datum $x \sim p_{\mathcal{X}}$. Notice that $\mathbb{E}_{\mathcal{X}}[\nabla_{\theta} \ell(x; \theta)] = \nabla_{\theta} \mathcal{L}(\theta)$, i.e., the SG is an unbiased estimator of the batch gradient $\nabla_{\theta} \mathcal{L}(\theta)$.

Crucially, the SG can be written in the form $\nabla_{\theta} \ell(x; \theta_t) = g(h(x; \theta_t); \theta_t)$. In other words, the back propagation only depends on the forward propagation through the context $h(x; \theta_t)$. The entire context must be kept in memory for computing the gradients. The context dominates the memory consumption in many applications.

ACT reduces the training memory footprint by compressing the context. Let $Q(h)$ be a compressor, which converts h

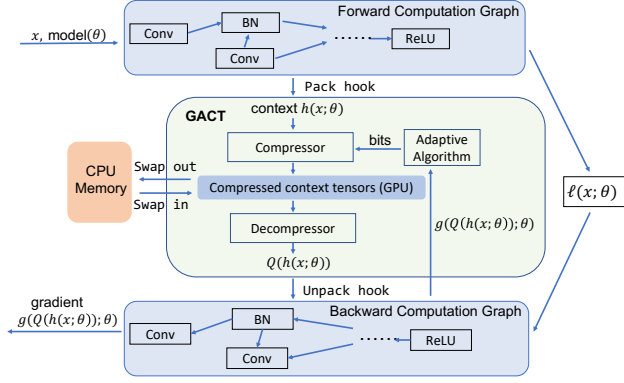


Figure 1. The architecture of GACT.

to compact formats while keeping $Q(h) \approx h$. Then, ACT computes the gradient with compressed context:

$$\theta_{t+1} \leftarrow \theta_t - \eta g(Q(h(x; \theta_t)); \theta_t). \quad (2)$$

We refer to $g(Q(h(x; \theta_t)); \theta_t)$ as the activation compressed (AC) gradient. ACT is significantly more memory efficient than the plain SGD, Eq. (1), since it only needs to store a compressed version of the context. Suppose the original context $h(x; \theta_t)$ consists of 32-bit floating point tensors, and $Q(\cdot)$ is a compressor which quantizes tensors to 2-bit integers, ACT will reduce the context memory by $16\times$. Fig. 1 illustrates the computational graph of ACT with these notations. In the following presentation, we might denote $h(x, \theta)$ simply by h when there is no confusion.

3.2. Convergence of ACT

ACT is a lossy approximation of SGD, as it uses an approximate gradient $g(Q(h); \theta)$. Therefore, some kind of theoretical guarantee is required for ACT to be useful. Fortunately, analyzing ACT is made significantly simpler by introducing an *unbiased stochastic* compressor $Q(\cdot)$, such that $\mathbb{E}_Q [Q(x)] = x$ for any x . $\mathbb{E}_Q [\cdot]$ means taking expectation over the compressor. In this way, $g(Q(h); \theta)$ can be viewed as a stochastic estimator of the batch gradient $\nabla \mathcal{L}(\theta)$, but the randomness comes not only from the datum x but also the compressor $Q(\cdot)$. Therefore, ACT is still an SGD algorithm. Standard analytical tools for SGD (Bottou et al., 2018) are applicable for studying ACT.

SGD algorithms have particular good properties when the SG is unbiased. In our case, this means $\mathbb{E}_Q [g(Q(h); \theta)] = g(h; \theta)$. However, the SG is biased general, even when the stochastic compressor itself is unbiased.¹

The key technique in this work is to construct an unbiased approximation of the AC gradient by linearizing the gradient function $g(\cdot; \theta)$. Consider the first-order Taylor expansion

¹Consider the example $g(h) = \mathbb{I}(h \geq 0.5)$, where $h \in [0, 1]$ and its AC gradient $g(Q(h)) = \mathbb{I}(Q(h) \geq 0.5)$ with the compressor $Q(h) \sim \text{Bernoulli}(h)$. Then, $\mathbb{E}[g(Q(h))] = P(Q(h) = 1) = h \neq g(h)$.

of $g(\cdot; \theta)$ at h :

$$\hat{g}(Q(h); h, \theta) := g(h; \theta) + J(h, \theta) \Delta h, \quad (3)$$

where $J(h, \theta) := \frac{\partial g(h; \theta)}{\partial h}$ is a Jacobian matrix, $\Delta h := Q(h) - h$ is the compression error. We further denote $\hat{g}_{x\theta}(Q(h); h) := \hat{g}(Q(h); h, \theta)|_{h=h(x; \theta)}$ and $J_{x\theta}(h) := J(h, \theta)|_{h=h(x; \theta)}$ for short. Since $\mathbb{E}[\Delta h(x; \theta)] = 0$, $\hat{g}_{x\theta}(Q(h); h)$ is an unbiased SG. Furthermore, the approximation error is small:

Proposition 1. *Assuming that $g(h; \theta)$ is twice differentiable w.r.t. h , and the second order derivative is bounded, then*

$$\mathbb{E}[\|g(Q(h); \theta) - \hat{g}_{x\theta}(Q(h); h)\|_2] = O(\text{Var}_Q[\Delta h]).$$

Since Δh itself is unbiased, $\text{Var}_Q[\Delta h] = \mathbb{E}_Q[\|\Delta h\|^2]$ is simply the expected compression error. Prop. 1 implies that the linearization error is bounded by the compression error. The linearized gradient \hat{g} is accurate if the compression is accurate. Using \hat{g} as a bridge, we arrive in the following convergence theorem:

Theorem 1. *Assume that:*

A1. $\mathcal{L}(\theta)$ is a continuous differentiable, $\nabla \mathcal{L}(\theta)$ is β -Lipschitz continuous.

A2. $\mathcal{L}(\theta)$ is bounded below by \mathcal{L}_* .

A3. $g(h; \theta)$ is differentiable w.r.t. h and $\exists b > 0$, s.t. $\forall \theta, \mathbb{E}[\|g(Q(h(x; \theta)); \theta) - \hat{g}_{x\theta}(Q(h); h)\|] \leq b$.

A4. $\exists \sigma^2 > 0$, s.t., $\forall \theta, \text{Var}[\hat{g}_{x\theta}(Q(h); h)] \leq \sigma^2$.

Then, for all $\eta < \frac{1}{2\beta}$, if we run ACT defined as Eq. (2) for T iterations, then we have

$$\min_{t=0, \dots, T-1} \mathbb{E}[\|\nabla \mathcal{L}(\theta_t)\|^2] \leq \frac{4(\mathcal{L}(\theta_0) - \mathcal{L}_*)}{\eta T} + 3b^2 + \eta\beta\sigma^2$$

Remark: The analytical technique used in Thm. 1 is rather standard, see Thm. 4.8 in Bottou et al. (2018). However, we consider the variance term σ^2 of the *linearized gradient*, rather than the SG itself. This formulation brings better analytical properties and an adaptive algorithm for determining the compression scheme, as we shall see soon in Sec. 4.

The convergence of ACT is affected by both the linearization error (A3) and the variance of the unbiased gradient $\hat{g}(\cdot; \theta)$ (A4). The latter is characterized as:

Proposition 2. $\text{Var}[\hat{g}_{x\theta}(Q(h); h)] = \text{Var}_{\mathcal{X}}[g(h; \theta)] + \mathbb{E}_{\mathcal{X}}[\text{Var}_Q[\hat{g}_{x\theta}(Q(h); h)]]$, where the second term on the RHS equals to $\mathbb{E}_{\mathcal{X}}[\text{Var}_Q[J_{x\theta}(h)\Delta h]] = O(\text{Var}_Q[\Delta h])$.

Prop. 2 separates the variance from different noise sources. $\text{Var}_{\mathcal{X}}[g(h(x, \theta); \theta)]$ is the variance raised by random sampling of data (“sampling variance”). $\mathbb{E}_{\mathcal{X}}[\text{Var}_Q[J_{x\theta}(h)\Delta h(x, \theta)]]$ is the variance raised by compression. Now, the convergence in Thm. 1 is depicted by $3b^2 + \eta\beta\sigma^2$. By Prop. 1, $b^2 = O(\text{Var}_Q[\Delta h]^2)$. By Prop. 2, $\sigma^2 = O(1) + O(\text{Var}_Q[\Delta h])$, since the sampling variance is not affected by compression. Therefore, when the compression is accurate ($\Delta h \rightarrow 0$), the impact of the linearization

error is negligible, and the variance of the unbiased gradient dominates. ACT behaves as if the AC gradient is unbiased.

4. Adapting the Compression Rate

In a network, some context tensors (such as those stored for computing the cross entropy loss) are extremely sensitive, a small amount of compression would result in diverged training, while other tensors are quite robust to compression. Therefore, we must apply different amounts of compression for each context tensor. As a general framework, we have no prior knowledge of the users' model architecture, so we designed an algorithm to infer the sensitivity for each context tensor and determine their compression rate automatically.

There is a tradeoff between the compression error and the storage requirement. We represent the storage requirement of the compressed context in *bits per dimension*. We assume that b_l bits/dim. are used for compression $h^{(l)}$, and $Q_{b_l}(h^{(l)})$ be the compression result. Let $b = (b_l)_{l=1}^L$ be a *compression scheme*, $Q_b(h) := \{Q_{b_l}(h^{(l)})\}_{l=1}^L$, and $\Delta_b h = Q_b(h) - h$.

4.1. Structure of Variance

As discussed in Sec. 3.2, when the compression is relatively accurate, the variance plays the main role in determining the convergence. Therefore, we would like to investigate how the compression scheme would impact the variance. Formally, we are interested in:

$$V(b; h, \theta) := \text{Var}_Q [\hat{g}(Q_b(h); h, \theta)].$$

Once $V(b; h, \theta)$ is known, we can find the minimum variance compression scheme under a given total bits budget B , by solving the integer programming problem:

$$\min_b V(b; h(x; \theta), \theta), \quad \text{s.t.} \quad \sum_{i=1}^L b_i D_i \leq B, \quad (4)$$

where D_l is the dimensionality of $h^{(l)}$. To proceed, we need the following assumptions on the compressor $Q_b(\cdot)$:

Assumption B1: The compressed result is element-wise uncorrelated. That is, for any $i \neq j$, $\text{Cov}[Q_b(h)_i, Q_b(h)_j] = 0$.

Assumption B2: For compressing $h^{(l)}(x; \theta)$ to b_l bits/dim., the compression error can be written in the form $\text{Var}[\Delta_{b_l} h^{(l)}(x; \theta)_j] \leq R_{lj}(x; \theta) S(b_l)$, where $S(b_l)$ is a known function. This isolates the effect of b_l through the unary factor $S(b_l)$.

Both assumptions can be achieved by a stochastic rounding quantizer (Courbariaux et al., 2015), where $R_{lj}(x; \theta) = \frac{1}{4} (\max_k h_k^{(l)} - \min_k h_k^{(l)})^2$ and $S(b) = (2^{b_l} - 1)^{-2}$. See Appendix A.4 for the derivations.

The following theorem reveals the structure of the variance:

Theorem 2. *Under assumptions B1, B2, there exists a family of functions $\{c_l(h, \theta)\}_{l=1}^L$, such that the compression*

Algorithm 1 Numerical algorithm for computing $c_l(h, \theta)$.

Require: A gradient evaluation function $g(\cdot; \theta)$

Require: A series of $L + 1$ random seeds $(r_l)_{l=1}^{L+1}$.

Require: Any compression scheme $b = (b_l)_{l=1}^L$

$\forall l$, seed $Q^{(l)}$ with r_l

$g_0 \leftarrow g(Q_b(h); \theta)$ {First iteration}

$\forall l$, seed $Q^{(l)}$ with r_l

seed $Q^{(l)}$ with r_{L+1}

$g_1 \leftarrow g(Q_b(h); \theta)$ {Second iteration, with another seed}

Return $\frac{1}{2} \|g_0 - g_1\|^2 / S(b_l)$

variance can be written in the form

$$V(b; h, \theta) \leq \sum_{l=1}^L c_l(h, \theta) S(b_l). \quad (5)$$

4.2. Computing Sensitivity

Thm. 2 reveals two good properties of the variance: (1) the impact of compressing different context tensors simply sums up, without affecting each other; and (2) the compression scheme only impacts the variance through $S(b_l)$. Both properties are brought about by linearization. Since $S(\cdot)$ is a known function, we only need to know $c_l(h, \theta)$ to solve problem Eq. (4). $c_l(h, \theta)$ can be understood as the sensitivity of the AC gradient to the compression of the l -th tensor. We can compute $c_l(h, \theta)$ numerically by leveraging the idempotence of compressing a tensor:

Assumption B3: If $h = Q(h')$ for some h' with non-zero probability, then $Q(h) = h$ and $\text{Var}_Q[Q(h)] = 0$.

Let $Q_b^{-l}(h) = \{Q_{b_1}(h^{(1)}), \dots, h^{(l)}, \dots, Q_{b_L}(h^{(L)})\}$ be some tensors, where every tensor except $h^{(l)}$ is compressed. Plug $h = Q_b^{-l}(h)$ into Eq. (5), and use B3, we have

$$V(b; Q_b^{-l}(h), \theta) \leq c_l(Q_b^{-l}(h), \theta) S(b_l).$$

The left hand side can be approximated by taking $\hat{g}(Q_b(h); h, \theta) \approx g(Q_b(h); \theta)$. Assume that $c_l(\cdot, \theta)$ is reasonably continuous, we have

$$c_l(h, \theta) \approx \text{Var}_Q [g(Q_b(h); \theta)]|_{h=Q_b^{-l}(h)} / S(b_l).$$

The variance can be replaced by empirical variance.

Alg. 1 illustrates this idea. To compute $\text{Var}_Q [g(Q_b(h); \theta)]$ at $h = Q_b^{-l}(h)$, we keep the random seeds fixed for all the compressors except the l -th one. We compute the empirical variance by two evaluations of $g(Q_b(h); \theta)$, which are two NN iterations (forward + backward propagation).

Finally, we assume that $c(h, \theta)$ remains stable for different mini-batches h , and along the training trajectory (θ_t) . Therefore, we maintain a c_l for each tensor l , which is updated by periodically running Alg. 1. Eq. (4) is approximately solved by the $O(L \log_2 L)$ greedy algorithm (Chen et al., 2021).

Another useful feature of this approach is predicting failure (in an *a posteriori* manner). If the compression vari-

```

1 import torch
2 import gact
3
4 model = ... # user defined model
5 controller = gact.controller(model, opt_level='L2')
6 controller.install_hook()
7
8 # training loop
9 for epoch in ...
10 for iter in ...
11     .....
12     # instruct gact how to perform forward and backward
13     def fwd_bwd_prop():
14         output = model(data)
15         loss = loss_func(output, target)
16         optimizer.zero_grad()
17         loss.backward()
18
19     controller.iterate(fwd_bwd_prop)

```

Figure 2. Usage example of GACT

ance $V(b; h, \theta)$ is dominating the overall gradient variance $\text{Var}[g(Q(h); \theta_t)]$, compression is adding too much noise to the gradient, and the convergence might be affected. The overall gradient variance can be computed by maintaining a running mean of the gradient. If $V(b; \theta)/\text{Var}[\hat{g}(Q(h); \theta_t)]$ is too large, we can raise an alert to the user to increase the storage budget.

5. System Implementation

We implemented GACT as a lightweight library in PyTorch. Users can use GACT for any NN architecture with several lines of code change. GACT uses low-level PyTorch hooks to capture context tensors, so it supports arbitrary operators, including custom operators defined by users. We implemented efficient CUDA kernels to infer tensor sensitivity and to perform compression during run time. GACT uses the same per-group quantizer in ActNN (Chen et al., 2021) as the compressor. However, GACT differs from ActNN in several aspects. ActNN relies on manual analytical deduction to compute the sensitivity for different operators, while GACT infers tensor sensitivity automatically, as described in Sec. 4.2. Moreover, ActNN performs layer-level quantization. It has to implement an activation compressed version for each operator and substitute operators during the training (e.g., replace `torch.nn.Conv2d` with `actnn.Conv2d`). In contrast, GACT runs at tensor level and uses a single hook interface to compress saved tensors for all operators.

5.1. General API

As shown in Fig. 2, the interface of GACT is straightforward and intuitive, requiring the user to (i) initialize the GACT controller and specify an optimization level (Line 5); (ii) install hooks (Line 6); and (iii) instruct GACT how to perform forward and backward propagation (Lines 13-17) and pass it as a function (`fwd_bwd_prop`) to the controller (Line 19). We require users to specify (iii) because GACT needs to numerically run the forward and backward pass to infer tensor sensitivity. Although `fwd_bwd_prop` is passed

to the controller every iteration, it is only called internally every `adapt_interval` iterations when tensor sensitivity changes. As shown in Sec. 6.1, tensor sensitivity stabilizes quickly after the first several epochs, `adapt_interval` can thus be set to a large number, introducing negligible impact on training speed.

5.2. System Architecture

Fig. 1 shows an overview of GACT. The GACT controller has three modules: Adaptive Algorithm; Compressor; and Decompressor. In the forward pass, the controller uses PyTorch `pack_hook` to capture all context tensors. Then Adaptive Algorithm infers tensor sensitivity based on gradients and assigns higher bits to more sensitive tensors, as described in Sec. 4.2. The bits information is used to instruct Compressor to perform quantization. In the backward pass, Decompressor dequantizes context tensors and uses `unpack_hook` to send the dequantized results back to the PyTorch’s auto differentiation engine. The controller is also responsible for swapping quantized tensors to the CPU and prefetching them back during the backward propagation if swapping is enabled.

5.3. Identifying Tensors to Quantize

The `pack_hook` and `unpack_hook` process all types of context tensors, including activation, parameters trained by the optimizer, and training states such as running mean/variance used by batch normalization. To guarantee that only the activations are quantized, we filter out saved parameters by recording the data pointers of all the model parameters before training, and we skip quantization if the input tensor pointer exists in the parameter pointer set. Similarly, GACT does not quantize training states by checking if the input tensor requires gradients.

However, using hooks blindly disables some memory-saving optimization. For example, in a transformer’s self-attention layer, the keys, query, value tensors are all calculated from the same input tensor. The saved objects of the three operations thus all refer to the same tensor. In this case, PyTorch triggers the `pack_hook` three times. If we perform quantization blindly, we waste computation resources and introduce extra memory consumption because the same underlying tensor is quantized and saved more than once. GACT avoids duplication by generating footprints for each input context tensor. We use the CUDA data pointer, sampled data points, and the tensor statistics (e.g., sum) as the footprint. GACT manages all quantized context tensors and uses the footprint to differentiate them. If a tensor is already quantized, GACT will skip quantization and return previous results directly.

5.4. Parallel Swap and Prefetch

To further reduce activation memory, we combine GACT with swapping. All compressed tensors are offloaded to the CPU during the forward pass and swapped back in the backward pass. Here, we replace the original tensor with quantized activation, as data movement is more expensive than computation. Swapping the original tensor saves the quantization overhead but adds more data movement cost between CPU and GPU. As shown in Sec. 6.4, quantization overhead is much smaller than copying full-precision data to CPU in modern GPU architecture.

Furthermore, we create two new streams (swap in/out) to parallelize the computation and swapping operation to reduce the swap overhead. The forward computation and swap-out process happen in parallel during the forward pass. During the backward pass, in each layer the swap-in stream is responsible for prefetching the compressed activation of the previous layer to avoid synchronization overhead. We leverage the CUDA event to ensure tasks in different streams are executed in the correct order.

5.5. Other Memory Optimizations

Gradient checkpointing. Gradient checkpointing (Chen et al., 2016) works by dividing the NN into segments. The algorithm only stores the inputs of each segment and recomputes the dropped activations segment by segment during backpropagation. The memory consumption is thus the cost of storing the inputs of all segments plus the maximum memory cost to backpropagate each segment. When combined with gradient checkpointing, GACT can reduce the memory consumption of both parts. GACT reduces the memory consumption of the first part by quantizing the segment inputs. Moreover, the activations saved during the recompute phase are also quantized, reducing the memory cost of the second part. Combining GACT with gradient checkpointing might introduce more training noise because the recompute starts from quantized segment inputs, making the forward pass of recompute phase not exact. However, in Sec. 6.4, we show the noise introduced by forwarding from the quantized tensors is negligible.

Memory efficient self-attention. When the batch size is very large, the single layer after dequantization occupies a large amount of memory and prevents the batch size from increasing further. We observe this problem in transformer-based models where self-attention has quadratic space complexity in terms of sequence length. To reduce the memory footprint of the self-attention layer, we implement the algorithm introduced in (Rabe & Staats, 2021) that achieves linear space complexity, and combines it with GACT.

Table 1. Optimization levels for GACT.

Level	Compression Strategy	Bits
L0	Do not compress	32
L1	per-group quantization with auto-precision	4
L2	L1 + swapping/prefetching	4
CB1	L1 + gradient checkpointing	4
CB2	CB1 + efficient self-attention	4

5.6. Optimization level

To exploit the trade-off between memory saving and training speed, GACT provides several optimization levels. Higher levels can save more memory but with more overhead. Tab. 1 lists these optimization levels. L1 uses per-group quantization with the adaptive algorithm. L2 combines per-group quantization with swapping and prefetching. For transformer-based models, CB1 combines GACT with gradient checkpointing. CB2 further reduces the peak memory by adding efficient self-attention to CB1.

6. Experiments

We first demonstrate the effectiveness of the GACT adaptive algorithm. We further apply GACT to a wide range of machine learning tasks, including image classification, object detection, text, and graph node classification. We compare the training accuracy and activation compression rate for full precision, adaptive 4/3/2 (using GACT to adaptively decide quantization bits with an average of 4/3/2 bit) and fix-4 bit (quantizing all tensors uniformly with 4 bits). Next, we study the trade-off between compression rate and training throughput and compare GACT with other state-of-the-art memory-saving methods. Lastly, we demonstrate the flexibility of GACT by exploring the possibility of combining it with other memory optimization methods (CB1, CB2 as listed in Table 1). We use open-source model implementations for all tasks.

6.1. Compression Strategy

We first test the effectiveness of our adaptive compression rate algorithm for training VGG-11 (Simonyan & Zisserman, 2015) on ImageNet. Fig. 3(a) plots the inferred per-tensor sensitivity c_l and the corresponding optimal bits/dim. GACT assigns more bits to more sensitive layers. The context tensor saved by the cross-entropy loss operator is most sensitive. A small amount of compression leads to a huge gradient variance. This makes sense since the loss is the first operator to back-propagate through, where the error accumulates. Therefore, GACT assigns 32 bits/dim. for the tensors in the classification head. With the adaptive algorithm, GACT with an average of 4 bits/dim. achieves smaller gradient variance than uniformly assigning 8 bits/dim. for all the tensors, as shown in Fig. 3(b). Finally, Fig. 3(c) shows that the sensitivity $c_l(h; \theta_t)$ remains stable during training. Therefore, periodically updating c_l at a large

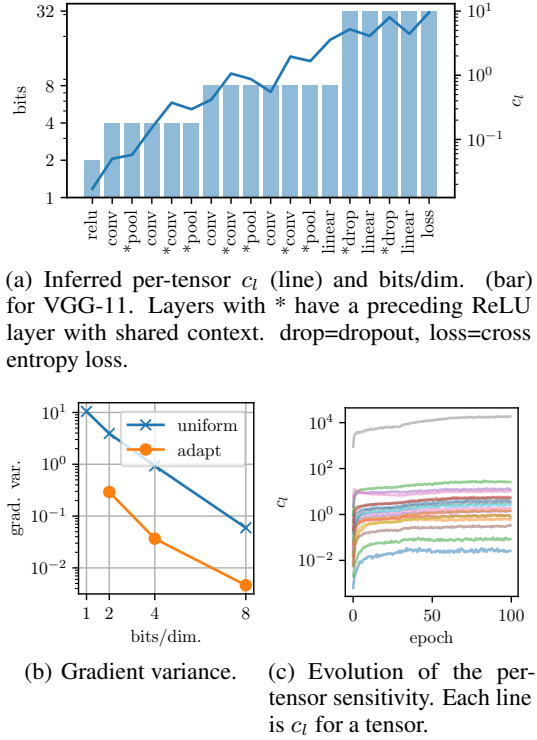


Figure 3. Effectiveness of the adaptive algorithm.

Table 2. For classification, we train VGG11 (Simonyan & Zisserman, 2015), ResNet-50 (He et al., 2016), and Swin-Tiny (Liu et al., 2021) on ImageNet (Deng et al., 2009). For object detection, we train RetinaNet (Lin et al., 2017), Faster R-CNN (Ren et al., 2015) on Coco (Lin et al., 2014). We report accuracy on validation sets (Div. indicates diverge) and the compression rate of context tensors (numbers in brackets) for both tasks.

Task	Model	FP32	GACT Adapt 4bit (L1)	GACT Adapt 2bit
Cls.	VGG11	68.75	68.77 (2.84×)	68.49 (3.34×)
	ResNet-50	77.29	76.96 (6.69×)	76.13 (11.39×)
	Swin-tiny	81.18	80.92 (7.44×)	77.91 (13.73×)
Det.	Faster RCNN	37.4	37.0 (4.86×)	36.1 (6.81 ×)
	RetinaNet	36.5	36.3 (3.11×)	Div.

interval is reasonable, and this introduces negligible impact on training speed.

6.2. Optimization level

We apply GACT on various computer vision tasks, including image classification and object detection, as shown in Fig. 2. We also vary the average bits used by the adaptive algorithm to explore the memory accuracy trade-off. On both tasks, GACT L1 achieves comparable (< 0.5% accuracy drop) or even better results than the full precision training, while reducing activation memory by up to 7.44×. Here, we list the accuracy of FP32 as the strongest accuracy baseline. For other lossy methods we consider in Sec. 6.3, the accuracy is no better than FP32, and we list their training accuracy in Appendix C. Notice that here GACT Adapt 2bit diverges

on the detection task. This is because, as shown in Sec.3.2, although ACT has unbiased gradients, the compression error and learning rate affect the convergence. When using 2 bit, the compression error is large and the learning rate has to be reduced accordingly to guarantee convergence. However, we do not want to slow training by decreasing the learning rate. All experiments are run with the same learning rate as the full precision. Therefore when compression error is large, the training diverges. Furthermore, we observe that the memory reduction varies among networks because GACT does not quantize intermediate states, and the size of intermediate states differs between networks. For example, in VGG11, when the batch size is 128, GACT reduces the saved tensor size from 5889MB to 2080MB, among which 78% (1494MB) is used to store the intermediate index for the max-pooling layer that is not quantized by GACT.

Next, we demonstrate the flexibility of GACT by applying it to a wider variety of natural language processing (NLP) and graph machine learning (Graph) tasks. We run multiple seeds for each task, and we report the mean±std of accuracy/F1 across runs as shown in Tab. 3. We include the detailed experimental setup in Appendix B. For both NLP and Graph tasks, GACT L1 achieves comparable training results with FP32, introducing less than 0.3% accuracy/F1-score drop, while reducing activation memory by 4.18× to 7.93×. Moreover, the results are stable across runs, introducing similar accuracy variance as FP32. We also show the training results of fix-4bit quantization, where all tensors are uniformly quantized with 4 bits. As shown in Tab. 3, fix-4 bit quantization causes significant accuracy/F1-score loss on various graph models. For Bert-large, fixed-4 bit quantization works fine because all the context tensors have similar sensitivity. On the other hand, GACT L1, using a similar amount of memory as always quantizing each layer to 4 bits, still performs on par with full precision training on all the models. This shows the necessity of using adaptive algorithms to assign bits based on tensor sensitivity for stabilized training. Moreover, for Bert-large and three graph models (GCN/GAT/GCNII), GACT converges and gives lossless results with 3 bits. Remarkably, across all the graph models, training with 2-bit GACT causes little accuracy loss (< 1%). This shows the robustness of our adaptive algorithm.

6.3. Memory Saving and Computational Overhead

Settings and baselines. We implement the benchmark with PyTorch 1.10 and measure the memory saving and overhead of GACT on an AWS g4dn.4xlarge instance, which has a 16GB NVIDIA T4 GPU and 64GB CPU memory. On ResNet-50, we compare with ActNN (Chen et al., 2021), a dedicated quantization framework for convolutional NNs, and DTR (Kirisame et al., 2020), a state-of-the-art rematerialization method for dynamic graphs. "swap" is a simple

Table 3. Accuracy and activation compression rate for NLP and Graph tasks. Accuracy that drops > 1% is in italic font.

Model	Dataset	FP32	Fix 4bit	GACT Adapt 4bit (L1)	GACT Adapt 3bit	GACT Adapt 2bit
GCN	Flickr	51.17 ± 0.19	50.93 ± 0.16 (7.56×)	51.08 ± 0.18 (7.93×)	51.14 ± 0.18 (11.34×)	51.20 ± 0.18 (17.56×)
	Reddit	95.33 ± 0.07	94.42 ± 0.11 (7.55×)	95.32 ± 0.07 (7.90×)	95.31 ± 0.07 (9.70×)	95.34 ± 0.06 (13.68×)
	Yelp	39.86 ± 0.94	39.85 ± 1.22 (5.94×)	40.06 ± 0.74 (6.42×)	40.21 ± 0.82 (7.46×)	39.89 ± 1.45 (9.00×)
	ogbn-arxiv	71.51 ± 0.65	68.61 ± 0.77 (7.54×)	71.35 ± 0.36 (8.09×)	70.82 ± 0.95 (10.45×)	70.87 ± 0.66 (13.75×)
GAT	Flickr	52.40 ± 0.28	35.24 ± 11.90 (4.23×)	52.26 ± 0.31 (4.34×)	51.68 ± 1.13 (5.04×)	51.62 ± 1.19 (5.46×)
	Reddit	95.95 ± 0.06	59.37 ± 11.48 (4.12×)	96.02 ± 0.09 (4.29×)	95.96 ± 0.06 (4.64×)	95.82 ± 0.06 (5.24×)
	Yelp	52.41 ± 0.69	36.09 ± 13.70 (4.04×)	52.18 ± 0.38 (4.18×)	51.63 ± 0.83 (4.53×)	51.15 ± 0.53 (5.24×)
	ogbn-arxiv	71.68 ± 0.54	54.64 ± 5.62 (5.04×)	71.80 ± 0.47 (5.09×)	71.47 ± 0.50 (6.14×)	71.21 ± 0.68 (6.98×)
GCNII	Flickr	52.37 ± 0.16	52.28 ± 0.16 (4.84×)	52.31 ± 0.16 (4.91×)	52.36 ± 0.16 (5.54×)	52.23 ± 0.15 (6.44×)
	Reddit	96.32 ± 0.24	86.50 ± 1.08 (4.51×)	96.11 ± 0.22 (4.52×)	96.01 ± 0.33 (5.16×)	95.54 ± 0.29 (5.92×)
	Yelp	62.33 ± 0.20	62.21 ± 0.22 (5.26×)	62.28 ± 0.26 (5.34×)	62.53 ± 0.36 (6.29×)	62.33 ± 0.37 (7.28×)
	ogbn-arxiv	72.52 ± 0.12	44.57 ± 5.01 (6.54×)	72.28 ± 0.35 (6.74×)	72.22 ± 0.28 (7.98×)	71.74 ± 0.26 (10.24×)
Bert-large	MNLI	86.74 ± 0.24	85.98 ± 0.16 (7.55×)	86.61 ± 0.11 (7.38×)	86.68 ± 0.08 (9.13×)	84.24 ± 0.74 (12.87×)
	SST-2	93.69 ± 0.30	93.46 ± 0.23 (7.55×)	93.54 ± 0.52 (7.30×)	93.20 ± 0.37 (9.05×)	91.90 ± 1.04 (12.91×)
	MRPC	88.20 ± 0.02	87.36 ± 0.19 (7.55×)	87.90 ± 0.10 (7.40×)	87.69 ± 0.07 (9.19×)	82.54 ± 0.38 (12.91×)
	QNLI	92.29 ± 0.14	92.34 ± 0.07 (7.55×)	92.44 ± 0.07 (7.42×)	92.43 ± 0.31 (9.19×)	90.74 ± 0.13 (12.95×)

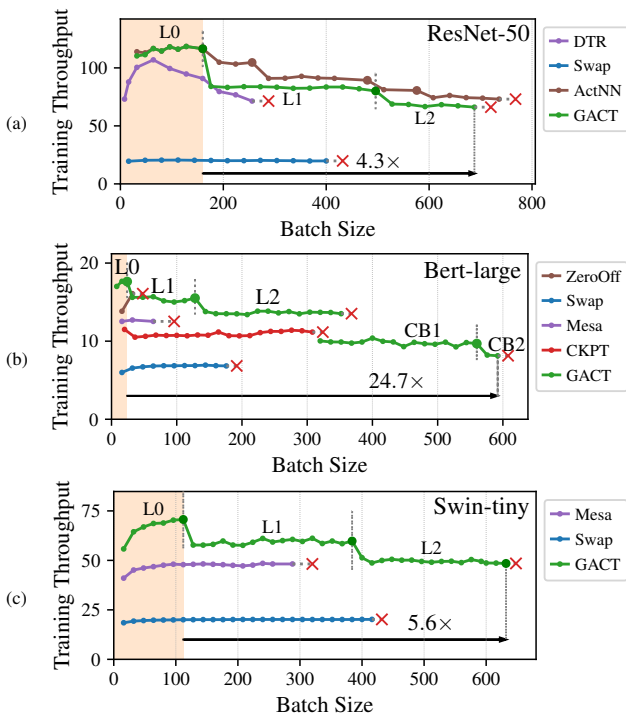


Figure 4. Training throughput vs batch size. Red cross mark means out-of-memory. The shaded yellow region denotes the batch sizes with full precision training given the memory budget. CKPT: Gradient checkpointing, ZeroOff: ZeRO-Offload.

swapping strategy that swaps all activations to the CPU. For Bert-large, we also show the results on Mesa (Pan et al., 2021), a memory-saving resource-efficient training framework for transformers, and ZeRO-Offload (Ren et al., 2021), a highly optimized system for training large-scale language models. Gradient checkpointing uses the default checkpointing policy provided by the transformer library (Wolf et al., 2020), where only the input to each transformer block is saved before the backward pass. On Swin-tiny, we only include Mesa and swap because other baselines lack the support for this network.

Results. We compare the training throughput of GACT against other memory saving systems in Fig. 4. On ResNet-50, GACT achieves similar throughput as ActNN (ActNN optimization L5 is not listed because it optimizes PyTorch memory allocation, which is unrelated to quantization and can also be applied to GACT), but ActNN enables training with a larger batch size. This is expected because ActNN implements efficient, customized layers for different operators in convolutional NNs. For Bert-large, Zero-offload fails quickly because it only offloads optimizer states that occupy a small portion of total memory to CPU. GACT L1 outperforms Mesa because Mesa only compresses tensors to 8 bit. When the batch is bigger, the activation size of each segment becomes the memory bottleneck and prevents gradient checkpointing from increasing the batch size. Moreover, combining GACT with gradient checkpointing and efficient self-attention further reduces the peak memory, increasing the batch size by up to 24.7×. Meanwhile, it introduces a small throughput overhead compared with the original gradient checkpointing. Across all the network architectures, GACT enables training with a 4.2× to 24.9× larger batch size under the same memory budget.

Network scaling. With GACT, we can construct larger models or train with a higher image resolution. Tab. 4 compares the largest model we can train against full precision. With the same batch size and memory budget, GACT can scale a ResNet-152 to 7.0× deeper, 3.6× wider or 3.0× higher resolution. Similarly, Bert-large can be scaled to 2.0× deeper or 1.6× wider. In GCN, GACT enables training 10.0× deeper and 1.7× wider network. Overall, GACT maintains 75% - 136% original training throughput.

6.4. Other Optimizations

We evaluate the idea of combining GACT with swapping on Bert-large-cased. As shown in Tab. 5, swapping compressed tensors is faster than swapping the original ones because communication between CPU and GPU is more time-consuming than computation. Combining GACT with

Table 4. Largest models GACT can train with 16G GPU memory. In ResNet (batch size=64), D (depth): number of layers, W (width): base width of the bottleneck block, R (resolution): width and height of input images. In Bert-large (batch size=16) and GCN, D (depth): number of transformer/gcn blocks, W (width): hidden size.

Dim	Maximum Value			Throughput (TFLOPS)			
	FP	L1	L2	FP	L1	L2	
ResNet-152	D	160	460	1124	0.43	0.47	0.41
	W	88	304	320	0.44	0.89	0.6
	R	232	548	716	0.41	0.39	0.44
Bert-large	D	32	56	64	0.67	0.56	0.53
	W	1280	1488	2032	0.68	0.61	0.60
GCN	D	24	152	240	0.20	0.14	0.15
	W	2464	3948	4244	0.36	0.38	0.40

Table 5. Swap and prefetch speed/memory on Bert-large.

Algorithm	Speed (sequence/s)	Peak Mem. (MB)	Total Mem. (MB)
FP32	16.41	9573	9527
FP32 + swap	6.02	5215	5093
GACT swap	12.95	5426	5325
GACT swap + prefetch	14.02	5426	5324

swapping increases training speed by up to 2.3×. Notice here that the peak memory use of “GACT swap” is slightly higher than “FP32 + swap” because GACT does not quantize and swap intermediate states such as running mean/var of BatchNorm layer. Moreover, prefetch increases the speed by about 7% with negligible memory overhead.

We next demonstrate combining GACT with gradient checkpointing (CB1). Gradient checkpointing is performed at the beginning of each transformer block, thus avoiding saving tensors generated within the block. We then apply GACT with gradient checkpointing, where the saved tensors are quantized with 4 bits. As shown in Tab. 6, the accuracy is unaffected. We also compare the activation memory and peak memory of CB1 and CB2 in Tab. 7. AM2 denotes the peak activation memory, which is the size of saved tensors after reforwarding the first transformer block. When batch size = 288, compared with gradient checkpointing on full precision (FP32), CB1 and CB2 reduce the peak activation size by 4.7× and 5.4× respectively.

7. Conclusion

This paper presents GACT, an ACT framework for generic NN architectures. We prove the convergence of GACT without prior knowledge about operator type or network architecture by analyzing a linearized approximation of ATC’s gradients. With the adaptive algorithm, GACT achieves neg-

Table 6. Accuracy of Bert-large-cased on SST-2 and QNLI datasets

Algorithm	SST-2	QNLI	Algorithm	SST-2	QNLI
FP32	93.58	92.42	CB1	93.81	92.26

Table 7. Memory use of different algorithms on Bert-large. AM1: Activation size before backward, AM2: Activation size after reforwarding the first transformer block. When batch size = 288, L0 runs out of memory, and therefore it is not listed below.

Batch Size	Algorithm	AM1 (MB)	AM2 (MB)	Peak Mem. (MB)
16	L0	4434	-	9573
	FP32 + CKPT	210	394	5541
	CB1	37	99	5286
	CB2	31	79	5269
288	FP32 + CKPT	3783	7092	12885
	CB1	515	1497	8251
	CB2	486	1307	8102

ligible accuracy loss on various tasks, reducing activation memory by up to 8.1× and enabling training with up to 24.7× batch size compared with full precision training.

Acknowledgements

This work was supported by the National Key Research and Development Project of China (No. 2021ZD0110502); NSF of China Project (No. 62106120), by the National Science Foundation through grants IIS-1955488, IIS-2027575, CCF-1723352, ARO W911NF2110339, ONR N00014-21-1-2724, and DOE award DE-SC0016260. We would also like to acknowledge partial support from DARPA, IARPA, the Sloan Foundation, NSF, and ONR. Our conclusions do not necessarily reflect the position or the policy of our sponsors, and no official endorsement should be inferred.

References

Anonymous. EXACT: Scalable graph neural networks training via extreme activation compression. In *Submitted to The Tenth International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=vkaMaq95_rX. under review.

Banner, R., Hubara, I., Hoffer, E., and Soudry, D. Scalable methods for 8-bit training of neural networks. In *Advances in Neural Information Processing Systems*, pp. 5145–5153, 2018.

Beaumont, O., Eyraud-Dubois, L., and Shilova, A. Efficient combination of rematerialization and offloading for training dnns. *Advances in Neural Information Processing Systems*, 34, 2021.

Bottou, L. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT’2010*, pp. 177–186. Springer, 2010.

Bottou, L., Curtis, F. E., and Nocedal, J. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.

Cen, Y., Hou, Z., Wang, Y., Chen, Q., Luo, Y., Yao, X., Zeng, A., Guo, S., Zhang, P., Dai, G., Wang, Y., Zhou, C.,

- Yang, H., and Tang, J. Cogdl: Toolkit for deep learning on graphs. *arXiv preprint arXiv:2103.00959*, 2021.
- Chakrabarti, A. and Moseley, B. Backprop with approximate activations for memory-efficient network training. *arXiv preprint arXiv:1901.07988*, 2019.
- Chen, J., Gai, Y., Yao, Z., Mahoney, M. W., and Gonzalez, J. E. A statistical framework for low-bitwidth training of deep neural networks. In *Advances in neural information processing systems*, 2020a.
- Chen, J., Zheng, L., Yao, Z., Wang, D., Stoica, I., Mahoney, M. W., and Gonzalez, J. E. Actnn: Reducing training memory footprint via 2-bit activation compressed training. In *International Conference on Machine Learning*, 2021.
- Chen, M., Wei, Z., Huang, Z., Ding, B., and Li, Y. Simple and deep graph convolutional networks. In *International Conference on Machine Learning*, pp. 1725–1735. PMLR, 2020b.
- Chen, T., Xu, B., Zhang, C., and Guestrin, C. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016.
- Courbariaux, M., Bengio, Y., and David, J.-P. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pp. 3123–3131, 2015.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In *2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171–4186, 2019.
- Evans, R. D. and Aamodt, T. AC-GC: Lossy activation compression with guaranteed convergence. *Advances in Neural Information Processing Systems*, 34, 2021.
- Evans, R. D., Liu, L., and Aamodt, T. M. Jpeg-act: accelerating deep learning via transform-based lossy compression. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pp. 860–873. IEEE, 2020.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.
- Fu, F., Hu, Y., He, Y., Jiang, J., Shao, Y., Zhang, C., and Cui, B. Don’t waste your bits! squeeze activations and gradients for deep neural networks via tinyscript. In *International Conference on Machine Learning*, pp. 3304–3314. PMLR, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B., Catasta, M., and Leskovec, J. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.
- Huang, C.-C., Jin, G., and Li, J. Swapadvisor: Pushing deep learning beyond the gpu memory limit via smart swapping. In *Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 1341–1355, 2020.
- Jain, P., Jain, A., Nrusimha, A., Gholami, A., Abbeel, P., Keutzer, K., Stoica, I., and Gonzalez, J. E. Checkmate: Breaking the memory wall with optimal tensor rematerialization. *arXiv preprint arXiv:1910.02653*, 2019.
- Jin, S., Li, G., Song, S. L., and Tao, D. A novel memory-efficient deep learning training framework via error-bounded lossy compression. In *26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pp. 485–487, 2021.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Kirisame, M., Lyubomirsky, S., Haan, A., Brennan, J., He, M., Roesch, J., Chen, T., and Tatlock, Z. Dynamic tensor rematerialization. *arXiv preprint arXiv:2006.09616*, 2020.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. Focal loss for dense object detection. In *International Conference on Computer Vision (ICCV)*, pp. 2980–2988, 2017.

- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. *International Conference on Computer Vision (ICCV)*, 2021.
- Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., and Wu, H. Mixed precision training. In *International Conference on Learning Representations*, 2018.
- Pan, Z., Chen, P., He, H., Liu, J., Cai, J., and Zhuang, B. Mesa: A memory-saving training framework for transformers. *arXiv preprint arXiv:2111.11124*, 2021.
- Peng, X., Shi, X., Dai, H., Jin, H., Ma, W., Xiong, Q., Yang, F., and Qian, X. Capuchin: Tensor-based gpu memory management for deep learning. In *Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 891–905, 2020.
- Rabe, M. N. and Staats, C. Self-attention does not need $o(n^2)$ memory. *arXiv preprint arXiv:2112.05682*, 2021.
- Ren, J., Rajbhandari, S., Aminabadi, R. Y., Ruwase, O., Yang, S., Zhang, M., Li, D., and He, Y. Zero-offload: Democratizing billion-scale model training. *arXiv preprint arXiv:2101.06840*, 2021.
- Ren, S., He, K., Girshick, R., and Sun, J. Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2015.
- Sun, X., Wang, N., Chen, C.-Y., Ni, J., Agrawal, A., Cui, X., Venkataramani, S., El Maghraoui, K., Srinivasan, V. V., and Gopalakrishnan, K. Ultra-low precision 4-bit training of deep neural networks. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio, P., and Bengio, Y. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355, 2018a.
- Wang, L., Ye, J., Zhao, Y., Wu, W., Li, A., Song, S. L., Xu, Z., and Kraska, T. Superneurons: Dynamic GPU memory management for training deep neural networks. In *23rd ACM SIGPLAN symposium on principles and practice of parallel programming*, pp. 41–53, 2018b.
- Wang, N., Choi, J., Brand, D., Chen, C.-Y., and Gopalakrishnan, K. Training deep neural networks with 8-bit floating point numbers. In *Advances in Neural Information Processing Systems*, pp. 7675–7684, 2018c.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. Transformers: State-of-the-art natural language processing. In *2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45, Online, October 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- Wu, S., Li, G., Chen, F., and Shi, L. Training and inference with integers in deep neural networks. In *International Conference on Learning Representations*, 2018.
- Zeng, H., Zhou, H., Srivastava, A., Kannan, R., and Prasanna, V. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931*, 2019.

A. Proof of Theorems

A.1. Theorem 1: Convergence of ACT

Assume that:

A1. $\mathcal{L}(\theta)$ is a continuous differentiable, $\nabla\mathcal{L}(\theta)$ is β -Lipschitz continuous. .

A2. $\mathcal{L}(\theta)$ is bounded below by \mathcal{L}_* .

A3. $g(h; \theta)$ is differentiable w.r.t. h and $\exists b > 0$, s.t. $\forall \theta, \mathbb{E} \|g(Q(h(x, \theta)); \theta) - \hat{g}(h(x, \theta); \theta)\| \leq b$.

A4. $\exists \sigma^2 > 0$, s.t., $\forall \theta, \text{Var} [\hat{g}(h(x, \theta))] \leq \sigma^2$.

Then, for all $\eta < \frac{1}{2\beta}$, if we run ACT defined as Eq. (2) for T iterations, then we have

$$\min_{t=0, \dots, T-1} \mathbb{E} \left[\|\nabla\mathcal{L}(\theta_t)\|^2 \right] \leq \frac{4(\mathcal{L}(\theta_0) - \mathcal{L}_*)}{\eta T} + 3b^2 + \eta\beta\sigma^2$$

Proof. Denote $m := \nabla_{\theta}\mathcal{L}(\theta_t)$, $\epsilon := \hat{g}(h(x, \theta_t); \theta_t) - m$, $d := g(Q(h(x, \theta_t)); \theta_t) - \hat{g}(h(x, \theta_t); \theta_t)$. Then, by A3 and A4, we have

$$\begin{aligned} \mathbb{E} [\epsilon] &= \mathbb{E} [g(h(x, \theta_t); \theta_t) - \nabla_{\theta}\mathcal{L}(\theta_t)] + \mathbb{E} [\langle J(x, \theta_t), \Delta Q(h(x, \theta_t)) \rangle] \\ &= \langle J(x, \theta_t), \mathbb{E} [\Delta Q(h(x, \theta_t))] \rangle = 0. \end{aligned} \quad (6)$$

$$\mathbb{E} \left[\|\epsilon\|^2 \right] = \|\mathbb{E} [\epsilon]\|^2 + \text{Var} [\epsilon] = \text{Var} [\hat{g}(h(x, \theta_t); \theta_t)] \leq \sigma^2. \quad (7)$$

$$\mathbb{E} [\|d\|] \leq b. \quad (8)$$

By the definitions, the ACT dynamics can be written as

$$\theta_{t+1} \leftarrow \theta_t - \eta(m + d + \epsilon).$$

By A1, we have

$$\mathcal{L}(\theta_{t+1}) \leq \mathcal{L}(\theta_t) - \eta \langle m, m + d + \epsilon \rangle + \frac{\beta\eta^2}{2} \|m + d + \epsilon\|^2. \quad (9)$$

By Eq.s (6,8)

$$\mathbb{E} [\langle m, m + d + \epsilon \rangle] \geq \|m\|^2 - \|m\| \|d\| + \langle m, \mathbb{E} [\epsilon] \rangle \geq \|m\|^2 - \|m\| b. \quad (10)$$

By Eq.s (6,7,8), and $\|x + y\|^2 \leq 2\|x\|^2 + 2\|y\|^2$,

$$\mathbb{E} \left[\|m + d + \epsilon\|^2 \right] = \mathbb{E} \left[\|m + d\|^2 \right] + \text{Var} [\epsilon] \leq 2\mathbb{E} [\|m\|^2] + 2\mathbb{E} [\|d\|^2] + \text{Var} [\epsilon] = 2\mathbb{E} [\|m\|^2] + 2b^2 + \sigma^2. \quad (11)$$

Taking expectation on both sides of Eq. (9), plug in Eq.s (10, 11), and use $\eta < \frac{1}{2\beta}$, we have

$$\begin{aligned} \mathbb{E} [\mathcal{L}(\theta_{t+1})] &\leq \mathcal{L}(\theta_t) - \eta(\|m\|^2 - \|m\| b) + \frac{\beta\eta^2}{2} (2\mathbb{E} [\|m\|^2] + 2b^2 + \sigma^2). \\ &= \mathcal{L}(\theta_t) - (\eta - \beta\eta^2) \|m\|^2 + \eta \|m\| b + \frac{\beta\eta^2}{2} (2b^2 + \sigma^2) \\ &= \mathcal{L}(\theta_t) - \frac{\eta}{2} \|m\|^2 + \eta \|m\| b + \frac{\beta\eta^2}{2} (2b^2 + \sigma^2). \end{aligned}$$

Completing the squares,

$$\mathbb{E} [\mathcal{L}(\theta_{t+1})] \leq \mathcal{L}(\theta_t) - \frac{\eta}{2} (\|m\| - b)^2 + \frac{\beta\eta^2}{2} (2b^2 + \sigma^2).$$

Take expectation on both sides and sum up for $t = 0, \dots, T-1$,

$$\mathbb{E} [\mathcal{L}(\theta_T)] - \mathcal{L}(\theta_0) \leq -\frac{\eta}{2} \sum_{t=0}^{T-1} \mathbb{E} (\|\nabla\mathcal{L}(\theta_t)\| - b)^2 + \frac{\beta\eta^2 T}{2} (2b^2 + \sigma^2).$$

Reorganize the terms,

$$\mathbb{E}_t \left[\mathbb{E} (\|\nabla\mathcal{L}(\theta_t)\| - b)^2 \right] \leq \frac{2(\mathcal{L}(\theta_0) - \mathcal{L}(\theta_T))}{\eta T} + \eta\beta(2b^2 + \sigma^2).$$

Let $t_* = \text{argmin}_t \mathbb{E} [\|\nabla\mathcal{L}(\theta_t)\|]$, and use A1, we have

$$\mathbb{E} (\|\nabla\mathcal{L}(\theta_{t_*})\| - b)^2 \leq \frac{2(\mathcal{L}(\theta_0) - \mathcal{L}_*)}{\eta T} + \eta\beta(2b^2 + \sigma^2).$$

Use $(a + b)^2 \leq 2a^2 + 2b^2$, we have

$$\mathbb{E} \left[\|\nabla\mathcal{L}(\theta_{t_*})\|^2 \right] \leq \frac{4(\mathcal{L}(\theta_0) - \mathcal{L}_*)}{\eta T} + (2\beta\eta + 2)b^2 + \eta\beta\sigma^2 \leq \frac{4(\mathcal{L}(\theta_0) - \mathcal{L}_*)}{\eta T} + 3b^2 + \eta\beta\sigma^2.$$

A.2. Proposition 1: The Linearization Error

Proof. Consider the gradient function $g(Q(h(x; \theta); \theta))$, whose output is a P -dimensional vector. Since it is twice differentiable, we construct the Taylor's expansion at $h(x; \theta)$ with Lagrange remainder:

$$\exists H_1, \dots, H_P, \text{ s.t., } \forall i, \quad g_i(Q(h(x; \theta); \theta)) = g_i(h(x, \theta); \theta) + J_i(x, \theta) \Delta h(x, \theta) + \Delta h(x, \theta)^\top H_i \Delta h(x, \theta),$$

where $J_i(h(x; \theta), \theta) := \frac{\partial g_i(h(x; \theta); \theta)}{\partial h}$. By the assumption, there exists $P > 0$, such that the linearization error is

$$\|g(Q(h(x; \theta); \theta)) - \hat{g}(h(x; \theta); h(x; \theta), \theta)\|_1 = \sum_{i=1}^P \Delta h(x, \theta)^\top H_i \Delta h(x, \theta) \leq \gamma P \|\Delta h(x, \theta)\|^2.$$

Taking expectation,

$$\begin{aligned} \mathbb{E} [\|g(Q(h(x; \theta); \theta)); h(x; \theta), \theta - \hat{g}(h(x; \theta); h(x; \theta), \theta)\|_2] &\leq \mathbb{E} [\|g(Q(h(x; \theta); \theta)); \theta - \hat{g}(h(x; \theta); h(x; \theta), \theta)\|_1] \\ &\leq \gamma P \text{Var} [\Delta h(x, \theta)] = O(\text{Var} [\Delta h(x, \theta)]). \end{aligned}$$
□

A.3. Proposition 2: The Order of the Variance

The following proposition is convenient for isolating the different noise sources.

Proposition A. (*Law of Total Variance*)

$$\text{Var} [X] = \mathbb{E} [\text{Var} [X | Y]] + \text{Var} [\mathbb{E} [X | Y]].$$

Proof. By definition

$$\text{Var} [\hat{g}(h(x; \theta_t); h(x; \theta), \theta_t)] = \text{Var} [g(h(x, \theta); \theta)] + \text{Var} [J(h(x; \theta), \theta) \Delta h(x, \theta)],$$

where $\text{Var} [g(h(x, \theta); \theta)]$ is the noise introduced by subsampling the data x . By law of total variance,

$$\text{Var} [J(h(x; \theta), \theta) \Delta h(x, \theta)] = \mathbb{E}_{\mathcal{X}} [\text{Var}_Q [J(h(x; \theta); \theta_t) \Delta h(x, \theta)]] + \underbrace{\text{Var}_{\mathcal{X}} [\mathbb{E}_Q [J(h(x; \theta); \theta_t) \Delta h(x, \theta)]]}_{=0},$$

where

$$\begin{aligned} \text{Var}_Q [J(h(x; \theta); \theta_t) \Delta h(x, \theta)] &= \mathbb{E}_Q [\|J(h(x; \theta); \theta_t) \Delta h(x, \theta)\|^2] \leq \mathbb{E}_Q [\|J(h(x; \theta); \theta_t)\|^2 \|\Delta h(x, \theta)\|^2] \\ &= \|J(h(x; \theta); \theta_t)\|^2 \mathbb{E}_Q [\|\Delta h(x, \theta)\|^2] = O(\text{Var} [\Delta h(x, \theta)]). \end{aligned}$$
□

A.4. Proposition 3: The Structure of the Variance

Before investigating the structure of $\text{Var}_Q [J(x; \theta_t) \Delta h(x, \theta)]$, let's do some recap: the parameter θ_t is a P -dimensional vector; the context difference $\Delta h(x, \theta)$ is a D -dimensional vector, and $J(x; \theta_t)$ is a $P \times D$ matrix. Recall that $\Delta h(x, \theta)$ is the concatenation of L -vectors, $\Delta h^{(l)}(x, \theta)$, and let $J^{(l)}(x, \theta) := \frac{\partial g}{\partial h^{(l)}} g((h^{(l)}(x; \theta))_{l=1}^L, \theta)$, which is a $P \times D_l$ matrix. Furthermore, let $h_j^{(l)}(x, \theta)$ be the j -th dimension, and $J_j^{(l)}(x, \theta)$ be its j -th column.

To proceed, we need to make the following assumptions to the compressor $Q(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}^D$:

B1: The compressed result is element-wise uncorrelated. That is, for any $i \neq j$, $\text{Cov} [Q(h)_i, Q(h)_j] = 0$.

B2: For compressing a vector h to b bits, the compression variance of each dimension can be written in the form $\text{Var} [Q(h)_j] \leq R_j(h) S(b)$, where $S(\cdot)$ is a known function.

Both assumptions can be achieved by a stochastic rounding (Courbariaux et al., 2015) quantizer, where

$$Q(h)_j = \begin{cases} T_{h,b}^{-1}(\lceil T_{h,b}(h_j) \rceil) & \text{w.p. } T_{h,b}(h_j) - \lfloor T_{h,b}(h_j) \rfloor \\ T_{h,b}^{-1}(\lfloor T_{h,b}(h_j) \rfloor) & \text{otherwise} \end{cases},$$

where $T_{h,b}(h_j) = (2^b - 1) \frac{h_j - \min_j h}{\max_j h - \min_j h}$. Since each dimension is quantized independently, B1 is met. Moreover,

$$\text{Var} [Q(h)_j] \leq \frac{1}{4} \left(\frac{\max_j h - \min_j h}{h_j - \min_j h} \right)^2 (2^b - 1)^{-2} = R_j(h) S(b),$$

where

$$R_j(h) = \frac{1}{4} \left(\frac{\max_j h - \min_j h}{(h_j - \min_j h)} \right)^2, \quad S(b) = (2^b - 1)^{-2}.$$

Proof. By definition,

$$J(h; \theta) \Delta h = \sum_{l=1}^L \sum_{j=1}^{D_l} J_j^{(l)}(h; \theta_l) \Delta h_j^{(l)}.$$

Using Assumption B1, we have

$$\begin{aligned} \text{Var}_Q [J(h; \theta) \Delta h] &= \mathbb{E}_Q \left[\left\| \sum_{l=1}^L \sum_{j=1}^{D_l} J_j^{(l)}(h; \theta_l) \Delta h_j^{(l)} \right\|^2 \right] \\ &= \sum_{l=1}^L \sum_{j=1}^{D_l} \mathbb{E}_Q \left[\left\| J_j^{(l)}(h; \theta_l) \Delta h_j^{(l)} \right\|^2 \right] \\ &= \sum_{l=1}^L \sum_{j=1}^{D_l} \left\| J_j^{(l)}(h; \theta_l) \right\|^2 \text{Var}_Q \left[\Delta h_j^{(l)} \right] \end{aligned}$$

Using Assumption B2, we have

$$\text{Var}_Q [J(h; \theta) \Delta h] \leq \sum_{l=1}^L \sum_{j=1}^{D_l} \left\| J_j^{(l)}(h; \theta_l) \right\|^2 R_l(h) S(b_l) = \sum_{l=1}^L c_l(h, \theta) S(b_l),$$

where $c_l(\theta, h) := R_l(h) \left\| J_j^{(l)}(h; \theta_l) \right\|_F^2$. □

B. Experiment Setup

B.1. Node classification task on graphs

We conduct experiments on four node classification datasets with standard splits, including Flickr, Reddit, Yelp from GraphSAINT (Zeng et al., 2019), and ogbn-arxiv from Open Graph Benchmark (OGB) (Hu et al., 2020). The four datasets cover extensive downstream applications with different scales. We use accuracy as the evaluation metric for multi-class classification and micro-F1 for multi-label classification. We run ten seeds (0 to 9) and report the average accuracy across runs.

We evaluate GACT on three representative GNN models, including GCN (Kipf & Welling, 2016), GAT (Veličković et al., 2017), and GCNII (Chen et al., 2020b) under the full-batch training setting. All three models are implemented by CogDL (Cen et al., 2021), a toolkit for graph neural networks.

B.2. Text classification task

We select four largest datasets, MNLI, QQP, SST-2, and QNLI, from the GLUE benchmark (Wang et al., 2018a). The four datasets cover different aspects of natural language understanding, including sentiment classification, natural language inference and paraphrase detection. We use the mainstream transformer implementation (Wolf et al., 2020) to train Bert-large (Devlin et al., 2019). We run three seeds (42, 43, 44) and report F1 for QQP, accuracy for the others.

C. Training Accuracy of Baselines

For all the baselines we compared in Sec. 6.3, only ActNN, Mesa, and ZeRO-Offload are lossy methods. All other methods are lossless and have the same training accuracy as FP32. For ResNet-50 on ImageNet, the training accuracy for FP32, GACT, ActNN L2, and ActNN L3 are 77.3, 77.0, 77.4, and 76.9. For Bert-Large on SST-2, the accuracy for FP32, GACT, Mesa, and ZeRO-Offload are 93.7, 93.5, 93.8, and 93.3. For Swin-tiny on ImageNet, the training accuracy for FP32, GACT, and Mesa are 81.2, 81.0, and 81.3 respectively.