

# Activations and Gradients Compression for Model-Parallel Training\*

Mikhail Rudakov<sup>1,2</sup>, Aleksandr Beznosikov<sup>1,2</sup>, Yaroslav Kholodov<sup>1</sup>, and Alexander Gasnikov<sup>1,2</sup>

<sup>1</sup> Innopolis University, Innopolis, Russia,

<sup>2</sup> Moscow Institute of Physics and Technology, Moscow, Russia

**Abstract.** Large neural networks require enormous computational clusters of machines. Model-parallel training, when the model architecture is partitioned sequentially between workers, is a popular approach for training modern models. Information compression can be applied to decrease workers communication time, as it is often a bottleneck in such systems. This work explores how simultaneous compression of activations and gradients in model-parallel distributed training setup affects convergence. We analyze compression methods such as quantization and TopK compression, and also experiment with error compensation techniques. Moreover, we employ TopK with AQ-SGD per-batch error feedback approach. We conduct experiments on image classification and language model fine-tuning tasks. Our findings demonstrate that gradients require milder compression rates than activations. We observe that  $K = 10\%$  is the lowest TopK compression level, which does not harm model convergence severely. Experiments also show that models trained with TopK perform well only when compression is also applied during inference. We find that error feedback techniques do not improve model-parallel training compared to plain compression, but allow model inference without compression with almost no quality drop. Finally, when applied with the AQ-SGD approach, TopK stronger than with  $K = 30\%$  worsens model performance significantly.

**Keywords:** distributed learning · model parallelism · activation compression · gradient compression · error feedback

## 1 Introduction

Neural networks have recently emerged as an essential field of computer science research. Hundreds of articles are published yearly on machine learning, including computer vision, natural language processing, reinforcement learning, and generative models. One of the reasons for such success is the tremendous growth of the model and dataset sizes. Recent examples include GPT-4 [24] and

---

\*The research of A. Beznosikov was supported by Russian Science Foundation (project No. 23-11-00229).

BLOOM [28], with up to billions of trainable parameters. Such models also require terabytes of data for training. Recent ROOTS text corpus [20], used for BLOOM training, is 1.6 TB large. Therefore, such models require a lot of CPU and GPU memory. Such computational resources are available only to a few large institutions and companies that can afford big servers with large network bandwidth. Nevertheless, smaller companies and research groups can also attempt to use modern architectures by pooling computational resources with each other over the Internet.

In any case, training large models is impossible without using several computational devices, and hence without using a distributed learning approach [35]. In distributed machine learning, data-parallel and model-parallel are two major parallelization paradigms that allow neural networks to be trained on several machines. In data parallelism, the training data is split between worker computers [18]. Each worker contains a full copy of the model and updates it during each training step according to averaged change from all the workers. Shallue et al. in the work [31] show that data parallelism does not negatively affect model quality with nearly the same number of training iterations. Data-parallel model training is widely used in large models training frameworks: Horovod [30], Pytorch Data-Parallel [21].

The second approach, model parallelism, or pipeline parallelism (MP/PP), is orthogonal. In model parallelism, the model itself is divided into blocks of layers, each stored on a separate machine. During training, workers are organized into a chain, where adjacent nodes communicate the layers activations in the forward pass and gradients in the backward pass. The model-parallel technique has been extensively used to train large models in systems like Megatron [32], Deepspeed [26], Petals [5]. Moreover, the model-parallel approach usually uses the pipelining technique. It allows distributed systems to use resources efficiently by splitting each batch into micro-batches, which are then pipelined on the block-divided model as in, for example, Gpipe [17], Xpipe [11], PipeDream [14].

However, slow communication becomes a problem in model parallelism [7]. During a forward pass, workers send vectors of layers activations. During a backward pass, vectors of respective gradients are sent to update model parameters. Communication time may be a bottleneck, leaving workers idle waiting for the data, especially when model-parallel training happens on machines located worldwide over slow network connection.

Research in the area of compression in model-parallel setup is present in the literature. Various quantization schemes are applied to activations [9,8,6] or gradients [33]. Sparsification techniques, which are popular in data-parallel case for gradients compression [34,3], have only a few results in the model-parallel case for activations compression [4,12]. Several works also utilize error compensation techniques in the model-parallel setting for activations [38] and gradients [33] for efficient convergence. Despite the work done on model-parallel compression, results tend to be independent and mostly empirical, obtained for activations or gradients separately.

In this work, we present experiments on model-parallel compression. Our focus is to study simultaneous compression of activations and gradients for various compression operators. Moreover, we experiment with existing error feedback approaches to test their applicability with sparsification compression.

The work is organized as follows. Section 2 describes experiment design and compression operators and techniques used for experiments with model parallelism. Section 3 presents and discusses results of model-parallel experiments in various compression scenarios. Section 4 reviews related works on activations and gradients compression in distributed learning. Finally, Section 5 concludes the work.

## 1.1 Contributions

**Simultaneous activation and gradient compression:** We conduct experiments on compressing both activations and gradients in model-parallel training and empirically evaluate convergence of models trained with quantization and TopK compressors, showing that:

1. In the case of quantization, gradients are more sensitive to compression than activations, requiring less compression for good convergence;
2. The strongest TopK compression level that does not harm the convergence is Top10%, and compression has to be applied during inference to produce comparable validation quality.

**Experiments with error feedback:** We apply various error compensation techniques to activations and gradients compression in the model-parallel setup, including EF and EF21, and show that:

1. Models trained with TopK compression and error feedback produce good validation results even with no compression applied;
2. When applied to both activations and gradients, error feedback techniques do not improve model convergence.

**AQ-SGD with TopK compression:** We evaluate how AQ-SGD [38] error feedback for activations works with TopK compression and conclude it does not improve model convergence compared to plain TopK compression.

## 2 Methodology

In this section, we describe the overall design of model-parallel compression experiments and characteristics of evaluated compression operators.

### 2.1 Experiment Design: Model-Parallel Training with Compression

In our experiments, we conduct training and fine-tuning of large neural networks using model parallelism approach with compression between pipeline stages.

Figure 1 presents this approach, where the whole model is partitioned on two devices with single communication stage in between (model parallelism degree is two). Adjacent workers in the pipeline exchange activations in the forward pass and gradients of these activations in the backward pass.




Fig. 1: Model parallel training example. Model parallelism degree is two, with one compression block. Activations are compressed in the forward pass, and gradients are compressed in the backward pass.

We use compression operators to send less data, which reduces potential communication overhead. In our experiments, we compress both activations and gradients to evaluate the limits at which compression can be used without compromising model convergence and quality.

Since our work focuses only on these two objectives, we do not conduct real model-parallel training on several machines. Instead, we integrate compression to the model directly, compressing activations and gradients between pipeline blocks, where communication happens in a real distributed system. This approach is equivalent to model-parallel training in terms of convergence analysis.

To reduce communication overhead, we apply various compression operators to activations and their respective gradients. We analyze convergence of quantization and TopK compression operators, and also use error feedback and AQ-SGD techniques with TopK compression.

## 2.2 Quantization

Quantization is a compression technique that involves mapping floating point numbers to a finite set of integer values. In our experiments, we utilize uniform  $k$ -bit quantization with scaling. Input vector with floating values is mapped to  $[0; 1]$  interval with min-max scaling and then quantized to  $k$  levels. Decompression transforms these values back to original scale. In our experiments, we evaluate quantization to 2, 4, 6, 8 bits for activations and gradients.

### 2.3 TopK compression

TopK compression is a sparsification method with proven convergence rates for data-parallel setup [3]. The TopK operator selects the largest K% values (in absolute value) from the input vector, setting other values to zero. This compression technique selects only the most critical data to send, which are supposed to be the largest activations or gradients.

We evaluate several levels of TopK compression (Top50%, Top30%, Top20%, Top10%, Top5%) for each task, compressing both activations and gradients.

### 2.4 Error Feedback

Error feedback (EF) is an error compensation technique widely used in data-parallel training with compression applied [29,23,27]. EF technique compensates errors introduced by compression by sending the accumulated error during next communication rounds. Such approach ensures that eventually all information contributes to the learning process.

The original version of EF, proposed by Seide et al. [29], saves compression error and adds it during the next communication round. Given vector  $X$  and compression operation  $\mathcal{C}$ , error is computed as  $e = X - \mathcal{C}(X)$ . During the next iteration, the compressed message is the input value added to the previous error, i.e.,  $\mathcal{C}(X + e)$ . Error is propagated to update model parameters on each iteration.

An advanced version of error feedback, EF21, was introduced by Richtarik et al. [27]. It features the compression of the current and previous parameter states' differences. As the parameter values tend to converge, the difference between two consecutive iterations' parameters converge to zero, and thus compressing such values would produce less communication error. Given activations vector  $X_{i+1}$  and previously sent value  $g_i$ , the communicated message is  $\mathcal{C}(X_{i+1} - g_i)$ , and buffer value is updated as  $g_{i+1} = g_i + \mathcal{C}(X_{i+1} - g_i)$ .

Based on ideas of error feedback techniques, we also introduce EF-mixed compression schema for TopK compression. The idea is to compress the largest  $K/2\%$  values from the input vector  $X$  and  $K/2\%$  of the current error buffer. After that, the input vector and buffer are added and sent, resulting in the same  $K\%$  non-zero values. This approach tries to get the largest and most important input values and the largest communication errors for activations and gradients.

In the experiments, we use global error buffer, meaning the accumulated error is added to the next batch.

### 2.5 AQ-SGD compression

AQ-SGD [38] is a technique designed to work with model parallelism using a variant of error feedback to compensate compression errors. Similarly to EF21, AQ-SGD compresses and communicates the change of activation values throughout the training process. However, the error buffer is individual for each batch in the training set, resulting in large memory footprint. The original work uses

quantization for forward and backward compression, with the AQ-SGD technique applied only for activations.

We use the AQ-SGD approach combined with TopK compression in our experiments. We evaluated whether AQ-SGD remains effective under biased TopK compression applied for activation and gradient, instead of the quantization in the original work. We test Top50%, Top30%, Top20%, and Top10% compression with AQ-SGD error feedback for activations and plain TopK compression for gradients.

### 3 Results

In this section, we present and evaluate the results of our experiments with various activation compression techniques in model-parallel neural network training. We present experiments with the ResNet-18 network and CIFAR-10 dataset. These experiments include quantization and TopK compression, usage of error feedback (EF) techniques, and AQ-SGD experiments. We also discuss experiments on fine-tuning GPT-2 on the Wikitext dataset with TopK compression. Our experiment code is available at GitHub<sup>‡</sup>.

#### 3.1 Training ResNet18 on CIFAR-10

The first series of experiments involved training ResNet18 [15] convolutional neural network on CIFAR-10 [19] dataset of labeled images. Image classification quality is measured with training accuracy; also train and test loss are reported.

Experiments were conducted on a single Tesla P100-PCIE-16GB GPU with model parallelism degree of 4. Training epochs was set to 100 epoch with a batch size of 100. SGD optimizer with momentum 0.9 and weight decay 5e-4 was used. The learning rate is controlled by a cosine annealing scheduler, with an initial value of 0.01 and  $T_{max} = 200$ .

Our implementation is based on the repository by kuangliu<sup>§</sup>, which compares various image classification models on the CIFAR-10 dataset.

As a baseline, training with no compression was conducted. It achieved 93.0% average test accuracy over five runs.

**Quantization** Figure 2 shows train loss and uncompressed and compressed accuracy on the test set for quantization experiments. Table 1 presents the best test accuracy achieved for each compression experiment.

Firstly, we observe that gradients are more sensitive to quantization than activations. Test accuracy comparable to uncompressed scenario is achieved only when gradients are quantized to at least 6 bits, while activations may be quantized to 2 or 4 bits. Such a result aligns with the results described in [38,4], where the weak performance of 2 and 4 bits gradients compression is reported.

<sup>‡</sup><https://github.com/Glemhel/ActivationsGradientsCompressionForMP>

<sup>§</sup><https://github.com/kuangliu/pytorch-cifar>

Table 1: Quantization Experiments Training Results on ResNet18 and CIFAR-10

Compression Mode	Test accuracy (%), compression off	Test accuracy (%), with compression
No compression	<b>93.00</b>	<b>93.00</b>
fw4-bw8	<b>93.10</b>	<b>92.95</b>
fw4-bw6	<b>91.93</b>	<b>91.76</b>
fw4-bw4	83.39	82.66
fw4-bw2	65.11	64.27
fw2-bw8	77.09	<b>92.05</b>
fw2-bw6	84.87	<b>91.59</b>
fw2-bw4	81.32	83.92

fw[A]-bw[B] means compressing activations to A bits, gradients to B bits. Each cell represents best result over 5 runs, each run for 100 epochs. Model-parallel degree is 4, with 3 compression operations used.

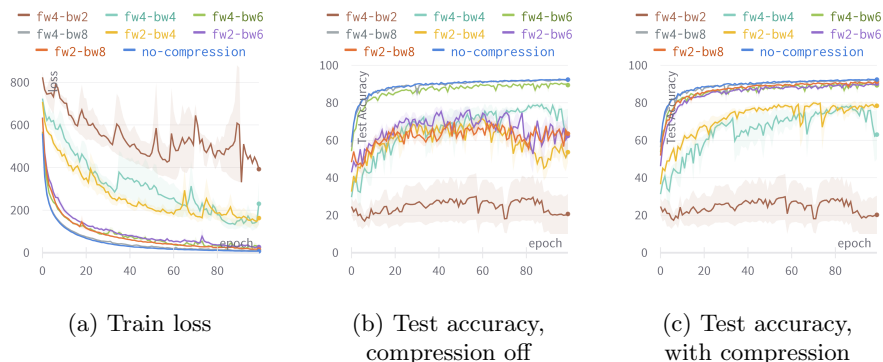


Fig. 2: Quantization experiments convergence on ResNet18 and CIFAR-10. fw[A]-bw[B] means compressing activations to A bits, gradients to B bits. Each line is average  $\pm$  standard error over 5 runs.

The distinction between compressing activations and gradients may lie in their impact. Compressing gradients can seriously disrupt optimization and convergence, while compressing activations has a milder effect, as long as the model can adapt to the compressed distribution.

Secondly, we report a notable difference in test accuracy with and without compression. Model test accuracy for the compressed case is 7-15 percentage points higher than inference without compression for quantization with activations compressed to 2 bits. We conclude that compression becomes part of the model because not using compression decreases model performance. When un-

compressed activations are passed into the model, they unpredictably change the model’s output.

Table 2: TopK Experiments Training Results on ResNet18 and CIFAR-10

Compression Mode	Test accuracy (%), compression off	Test accuracy (%), with compression
No compression	<b>93.00</b>	<b>93.00</b>
Top 50%	<b>93.16</b>	<b>93.38</b>
Top 30%	<b>91.26</b>	<b>93.06</b>
Top 20%	86.74	<b>92.73</b>
Top 10%	75.89	<b>91.87</b>
Top 5%	55.09	90.01
Top 2%	49.95	85.52

Each cell represents best result over 5 runs, each run for 100 epochs. Model-parallel degree is 4, with 3 compression operations used. Activations and gradients are compressed independently.




Fig. 3: TopK experiments convergence on ResNet18 and CIFAR-10. Each line is average  $\pm$  standard error over 5 runs. Model-parallel degree is 4, with 3 compression operations used. Activations and gradients are compressed independently.

**TopK Compression** Results of experiments with TopK compression are provided in Table 2, with learning curves displayed in Figure 3. Similarly to the



quantization case, we observe that inference with compression results in better test accuracy than without compression for TopK compression. While compressed inference test accuracy stays satisfactory up to Top10% compression, uncompressed test accuracy is below 90% already for Top20%. In the TopK compression case, activations that are usually set to 0 during compression change the model behavior significantly when no compression is applied. Therefore, for a model trained with compression, the application of compression during inference is obligatory to achieve good results. To the best of our knowledge, no previous works considered such aspect of model-parallel communication compression. Moreover, our results demonstrate that Top10% compression is a good technique for application in convolutional neural networks such as ResNet.

Table 3: Error Feedback Experiments Training Results on ResNet18 and CIFAR-10

Compression Mode	Test accuracy(%), compression off	Test accuracy(%), with compression
No compression	<b>93.00</b>	<b>93.00</b>
EF + Top 10%, warmup 20	<b>91.02</b>	89.85
EFmixed + Top 10%, warmup 20	61.37	90.97
EF21 + Top 5%	89.37	89.29
EF21 + Top 10%	90.83	<b>91.19</b>
EF21 + Top 10%, warmup 20	90.71	<b>91.77</b>

Each cell is a single run for 100 epochs.

Warmup 20 runs use uncompressed baseline weights after 20 epochs.

Model-parallel degree is 4, with 3 compression operations used.

Activations and gradients are compressed independently, with global EF batch buffer for each compression operator.

**Compression with Error Feedback** Table 3 and Figure 4 shows training results for error feedback experiments. Results suggest that using EF in addition to TopK compression does not improve the model accuracy. Significant difference of activations and gradients between batches may be why the EF-based approaches does not improve model quality when applied directly. Thus, adding errors to an example of a different target class may significantly affect the error in the buffer. This is probably why the global EF buffer fails to improve model performance in a model-parallel setup.

The unexpected yet positive finding is that the EF technique makes test accuracy results similar for uncompressed and compressed inference. Test accuracy without compression is only 1-2 percentage points below the one of compressed




Fig. 4: Error feedback experiments convergence on ResNet18 and CIFAR-10. Model-parallel degree is 4, with 3 compression operations used. Activations and gradients are compressed independently, with global EF batch buffer for each compression operator. Runs with suffix *base20* use uncompressed baseline weights after 20 epochs.

variant of EF or EF21, while for the plain TopK compression, the drop in test accuracy is at least 7-15 percentage points. We hypothesize that EF buffer is considered as noise to the model, and the model ignores such noise to the activations, becoming more robust and leading to comparable uncompressed and compressed test results.

Table 4: AQ-SGD and TopK Experiments Training Results on ResNet18 and CIFAR-10

Compression Mode	Test accuracy(%), compression off	Test accuracy(%), with compression
No compression	<b>93.00</b>	<b>93.00</b>
AQ-SGD + Top 50%, warmup 10	<b>92.44</b>	<b>92.54</b>
AQ-SGD + Top 30%, warmup 10	<b>91.86</b>	<b>91.61</b>
AQ-SGD + Top 20%, warmup 10	90.82	87.8
AQ-SGD + Top 10%, warmup 10	85.91	84.16

Each cell is a single run for 100 epochs.

Warmup 10 runs use uncompressed baseline weights after 10 epochs.

Model-parallel degree is 4, with 3 compression operations used.

Activations and gradients are compressed independently, with AQ-SGD per-example buffer applied only for activations.

**AQ-SGD with TopK compression** Finally, Table 4 presents results of experiments based on AQ-SGD [38] approach with TopK compression used instead of quantization.

Learning curves in Figure 5 suggest that convergence of the AQ-SGD method with TopK compression is not improved compared to the plain TopK compression. Test accuracy growth does not match the uncompressed baseline for Top10% compression. We suppose that convergence for TopK compression used with the AQ-SGD approach is slow due to the biasedness of TopK compression.

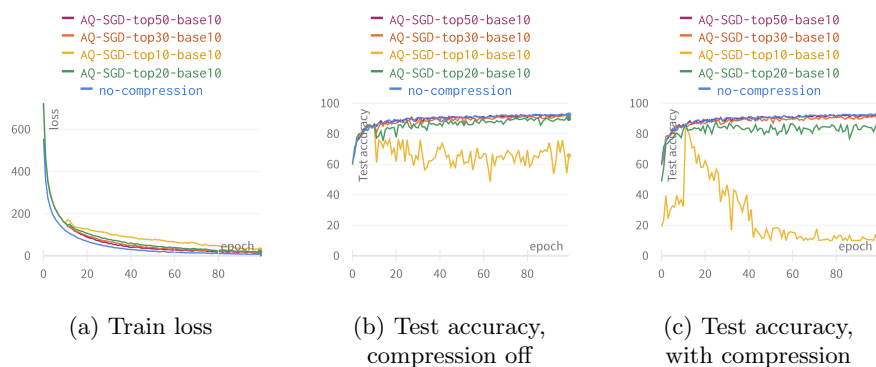


Fig. 5: AQ-SGD and TopK experiments convergence on ResNet18 and CIFAR-10. Model-parallel degree is 4, with 3 compression operations used.

Activations and gradients are compressed independently, with AQ-SGD per-example buffer applied only for activations. Runs with suffix *base10* use uncompressed baseline weights after 10 epochs.

### 3.2 Fine-tuning GPT-2 on Wikitext

We fine-tuned pretrained GPT-2-small [25] on Wikitext text dataset [22], wikitext-2-raw-v1 version. Quality was measured by test loss and perplexity metrics. We fine-tuned the model for 4 epochs with a batch size of 8. We used Hugging Face model fine-tuning code<sup>¶</sup> and integrated compression into it. Experiments were conducted with Tesla P100-PCIE-16GB GPU. Model parallelism degree was set to four.

Figure 6 shows train and eval loss during fine-tuning. Table 5 presents fine-tuning results with TopK compression applied. The first observation is that TopK compression starts to increase validation loss significantly already from Top20% compression. Compared to our experiments with ResNet18, where Top10% compression resulted in satisfactory results, current setup does not tolerate such

<sup>¶</sup>[https://github.com/huggingface/transformers/blob/main/examples/pytorch/language-modeling/run\\_clm.py](https://github.com/huggingface/transformers/blob/main/examples/pytorch/language-modeling/run_clm.py)

Table 5: TopK Compression Fine-tuning Experiments Results on GPT-2 and Wikitext

Compression Mode	Eval loss	Perplexity
No compression	<b>3.05</b>	<b>21.01</b>
Top 50%	<b>3.11</b>	<b>22.38</b>
Top 30%	<b>3.27</b>	<b>26.28</b>
Top 20%	3.51	33.32
Top 10%	4.31	74.51
Top 10% separate	8.0	2990.16

Model-parallel degree is 4, with 3 compression operations used. TopK compression reuses TopK indices from activations to compress gradients.

TopK separate mode compresses activations and gradients independently.

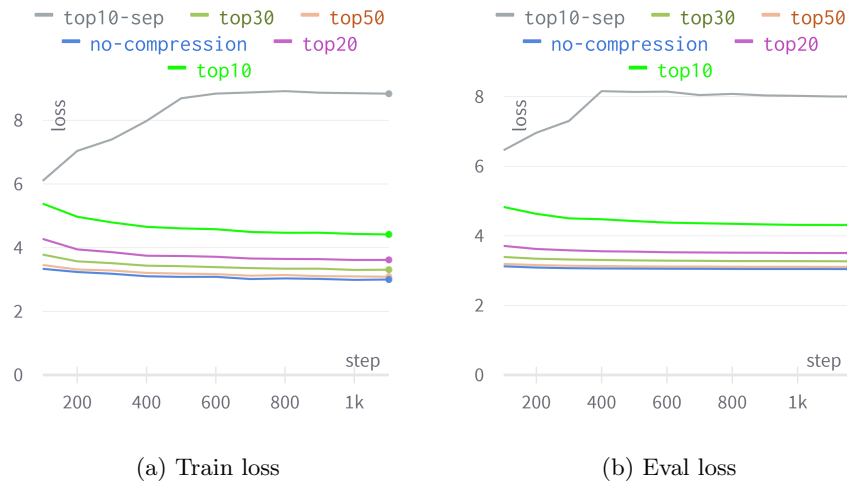


Fig. 6: TopK experiments convergence on GPT-2 and Wikitext.

compression. Two factors may explain this behaviour: the architectural differences between transformer and CNN models, and the use of fine-tuning instead of training from scratch, which may hinder efficient learning of TopK compression distributions.

The second observation is regarding the independence of activations and gradients compression. In the fine-tuning scenario, compressing activations and gradients independently leads to very large eval loss and even divergence of the model. We hypothesize that the fact we use pretrained GPT-2 defines such behavior. For the pre-trained model, zeroing out a large ratio of activations affects gradients more than the actual training sample model is fine-tuned on.

In our opinion, to prevent such behavior, reusing TopK% indices is necessary to achieve good results.

## 4 Related Works

This section describes related works on compression in distributed training. In general, compression techniques are used to reduce communication overhead, and can be applied to layers activations or their respective gradients.

### 4.1 Gradients compression

Gradients compression is naturally applied in the data-parallel approach. Several methods have been developed to reduce communication overhead without significantly decreasing model quality.

The most common way to reduce the space taken by floating point numbers is the quantization approach. Several works on robust quantization show that under certain assumptions and techniques used, quantized learning leads to the same convergence and model quality [29,2]. Alistarh et al. [1] developed a framework for evaluating convergence of quantization-based gradient compression methods. Quantization methods with distribution-aware choice of quantization levels achieve even better practical results applied to large models training [13,16,9]. Quantization methods have the advantage of being unbiased, hence being able to converge as with uncompressed training.

Another class of compression methods is based on low-rank approximation of gradients vector. Wang et al. [37] propose compression by sparsification of singular values of gradients. PowerSGD by Vogels et al. [36] compresses gradients even more with the concept of power iteration. Such methods also apply to model parallelism: Optimus-CC framework shows superior training time results with PowerSGD-based gradients compression [33]. However, the decomposition of gradients requires additional computation time for compression.

To cope with errors introduced by gradient compression, error feedback (EF) techniques have been proposed to improve convergence. Initially proposed by Seide et al. [29] as a heuristic, EF has shown to be an effective method for improving convergence for data-parallel tasks. The newer EF21 approach improves EF by communicating changes of gradients [27]. MARINA [10] and DIANA [23] use compression of gradient differences to lower convergence bounds for strongly convex and non-convex objectives in data-parallel setup. Optimus-CC [33] also compresses differences of gradients between mini-batches, resembling the EF21 approach, but applied to model parallelism. Overall, error feedback techniques achieve better theoretical and practical results in data parallel setting, but only a few works use EF with model parallelism.

Finally, various sparsification methods are also used for gradient compression. Beznosikov et al. [3] analyze the theoretical bounds of biased methods. They also propose improved TopK compression with exponential dithering to achieve more economy in communication with comparable model convergence in

a data-parallel setup. However, sparsification (for example, TopK compression) requires to send not only the selected data, but also the indices, which increases communication cost.

## 4.2 Activations compression

Similarly to gradient compression, activation compression could also be applied in both data- and model-parallel setups. Usually, it is used to reduce memory footprint on inference, as well as communication time in the case of model parallelism.

As the sizes of activations can be significant, some works compress activations even without model parallelism. Compressed activations are used to compute model gradients during a backward pass. AC-GC [8] uses fixpoint quantization to compress activations up to 15 times with only 0.1% average accuracy loss. Fu et al. [9] also use distribution-aware quantization to compress activations and gradients, decreasing memory consumption up to 2-4 times. A novel approach LLM-int8() [6] uses 8-bit quantization for speeding up inference of large language models without quality degradation. This approach also considers outliers values passed in the original fp32 format.

A notable application of the error feedback approach to activations compression has been introduced in AQ-SGD model-parallel compression framework [38]. This method communicates quantized differences of activations between pipeline blocks, achieving up to 8.5 times increase of total training throughput over slow networks.

Experimental work on activations compression methods by Bian et al. [4] compares popular compressors for model-parallel training and draws several conclusions on the practical application of activations compression. In particular, they observe that learning-based autoencoder compression performs best in terms of convergence and training throughput compared to quantization and TopK compression.

## 5 Conclusion

In this work, we conducted experiments on communication compression of activations and gradients for model-parallel neural networks distributed training. In particular, we explored how quantization, TopK compression, and error feedback-based approaches perform in several machine learning tasks.

We observed that model gradients are more sensitive to compression than activations as shown in quantization experiments. We also empirically showed that TopK compression for both activation and gradients can be applied with at most Top 10% to achieve comparable model quality. We evaluated error feedback techniques and noticed no significant improvement in convergence with TopK compression of activations and gradients in the model-parallel setup. However, error feedback techniques helped to overcome inference performance drop without compression applied. Finally, we reported that the AQ-SGD [38] approach

could not be applied with TopK compression stronger than  $K = 30\%$  to achieve satisfactory convergence.

Several potential limitations of our research should be acknowledged. Firstly, since model parallelism is usually applied for large models, this work lacks experiments with training state-of-the-art large models. Secondly, we test only one model parallelism pipeline and hyperparameters configuration for each experiment, which should be extended in the future work.

We propose several research directions based on our work: experiments with large language models, reducing AQ-SGD approach memory footprint, and explore more biased compression techniques apart from TopK.

## Acknowledgements

We deeply thank the most irreplaceable and the most secretive colleague from Yandex.Research for insightful discussions during our research. We also thank Anton Antonov for providing computational resources on early stages of the project.

## References

1. Alistarh, D., Grubic, D., Li, J., Tomioka, R., Vojnovic, M.: Qsgd: Communication-efficient sgd via gradient quantization and encoding. *Advances in neural information processing systems* **30** (2017)
2. Bernstein, J., Wang, Y.X., Azizzadenesheli, K., Anandkumar, A.: signsgd: Compressed optimisation for non-convex problems. In: *International Conference on Machine Learning*. pp. 560–569. PMLR (2018)
3. Beznosikov, A., Horváth, S., Richtárik, P., Safaryan, M.: On biased compression for distributed learning. *arXiv preprint arXiv:2002.12410* (2020)
4. Bian, S., Li, D., Wang, H., Xing, E.P., Venkataraman, S.: Does compressing activations help model parallel training? (2023)
5. Borzunov, A., Baranchuk, D., Dettmers, T., Ryabinin, M., Belkada, Y., Chumachenko, A., Samygin, P., Raffel, C.: Petals: Collaborative inference and fine-tuning of large models. *arXiv preprint arXiv:2209.01188* (2022)
6. Dettmers, T., Lewis, M., Belkada, Y., Zettlemoyer, L.: Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems* **35**, 30318–30332 (2022)
7. Diskin, M., Bukhtiyarov, A., Ryabinin, M., Saulnier, L., Sinitsin, A., Popov, D., Pyrkin, D.V., Kashirin, M., Borzunov, A., Villanova del Moral, A., et al.: Distributed deep learning in open collaborations. *Advances in Neural Information Processing Systems* **34**, 7879–7897 (2021)
8. Evans, R.D., Aamodt, T.: Ac-gc: Lossy activation compression with guaranteed convergence. *Advances in Neural Information Processing Systems* **34**, 27434–27448 (2021)
9. Fu, F., Hu, Y., He, Y., Jiang, J., Shao, Y., Zhang, C., 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)

10. Gorbunov, E., Burlachenko, K.P., Li, Z., Richtárik, P.: Marina: Faster non-convex distributed learning with compression. In: International Conference on Machine Learning. pp. 3788–3798. PMLR (2021)
11. Guan, L., Yin, W., Li, D., Lu, X.: Xpipe: Efficient pipeline model parallelism for multi-gpu dnn training. arXiv preprint arXiv:1911.04610 (2019)
12. Gupta, V., Choudhary, D., Tang, P., Wei, X., Wang, X., Huang, Y., Kejariwal, A., Ramchandran, K., Mahoney, M.W.: Training recommender systems at scale: Communication-efficient model and data parallelism. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining. pp. 2928–2936 (2021)
13. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. arXiv preprint arXiv:1510.00149 (2015)
14. Harlap, A., Narayanan, D., Phanishayee, A., Seshadri, V., Devanur, N., Ganger, G., Gibbons, P.: Pipedream: Fast and efficient pipeline parallel dnn training. arXiv preprint arXiv:1806.03377 (2018)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
16. Hong, C., Kim, H., Baik, S., Oh, J., Lee, K.M.: Daq: Channel-wise distribution-aware quantization for deep image super-resolution networks. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2675–2684 (2022)
17. Huang, Y., Cheng, Y., Bapna, A., Firat, O., Chen, D., Chen, M., Lee, H., Ngiam, J., Le, Q.V., Wu, Y., et al.: Gpipe: Efficient training of giant neural networks using pipeline parallelism. *Advances in neural information processing systems* **32** (2019)
18. Krizhevsky, A.: One weird trick for parallelizing convolutional neural networks. arXiv preprint arXiv:1404.5997 (2014)
19. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
20. Laurençon, H., Saulnier, L., Wang, T., Akiki, C., del Moral, A.V., Le Scao, T., Von Werra, L., Mou, C., Ponferrada, E.G., Nguyen, H., et al.: The bigscience roots corpus: A 1.6 tb composite multilingual dataset. In: Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (2022)
21. Li, S., Zhao, Y., Varma, R., Salpekar, O., Noordhuis, P., Li, T., Paszke, A., Smith, J., Vaughan, B., Damania, P., et al.: Pytorch distributed: Experiences on accelerating data parallel training. arXiv preprint arXiv:2006.15704 (2020)
22. Merity, S., Xiong, C., Bradbury, J., Socher, R.: Pointer sentinel mixture models. In: International Conference on Learning Representations (2016)
23. Mishchenko, K., Gorbunov, E., Takáč, M., Richtárik, P.: Distributed learning with compressed gradient differences. arXiv preprint arXiv:1901.09269 (2019)
24. OpenAI: Gpt-4 technical report (2023)
25. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI blog* **1**(8), 9 (2019)
26. Rasley, J., Rajbhandari, S., Ruwase, O., He, Y.: Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 3505–3506 (2020)
27. Richtárik, P., Sokolov, I., Fatkhullin, I.: Ef21: A new, simpler, theoretically better, and practically faster error feedback. *Advances in Neural Information Processing Systems* **34**, 4384–4396 (2021)



28. Scao, T.L., Fan, A., Akiki, C., Pavlick, E., Ilić, S., Hesslow, D., Castagné, R., Luccioni, A.S., Yvon, F., Gallé, M., et al.: Bloom: A 176b-parameter open-access multilingual language model. arXiv preprint arXiv:2211.05100 (2022)
29. Seide, F., Fu, H., Droppo, J., Li, G., Yu, D.: 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In: Fifteenth annual conference of the international speech communication association (2014)
30. Sergeev, A., Del Balso, M.: Horovod: fast and easy distributed deep learning in tensorflow. arXiv preprint arXiv:1802.05799 (2018)
31. Shallue, C.J., Lee, J., Antognini, J., Sohl-Dickstein, J., Frostig, R., Dahl, G.E.: Measuring the effects of data parallelism on neural network training. arXiv preprint arXiv:1811.03600 (2018)
32. Shoyebi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., Catanzaro, B.: Megatron-lm: Training multi-billion parameter language models using model parallelism. arXiv preprint arXiv:1909.08053 (2019)
33. Song, J., Yim, J., Jung, J., Jang, H., Kim, H.J., Kim, Y., Lee, J.: Optimus-cc: Efficient large nlp model training with 3d parallelism aware communication compression. In: Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2. pp. 560–573 (2023)
34. Stich, S.U., Cordonnier, J.B., Jaggi, M.: Sparsified sgd with memory. *Advances in Neural Information Processing Systems* **31** (2018)
35. Verbraeken, J., Wolting, M., Katzy, J., Kloppenburg, J., Verbelen, T., Rellermeyer, J.S.: A survey on distributed machine learning. *Acm computing surveys (csur)* **53**(2), 1–33 (2020)
36. Vogels, T., Karimireddy, S.P., Jaggi, M.: Powersgd: Practical low-rank gradient compression for distributed optimization. *Advances in Neural Information Processing Systems* **32** (2019)
37. Wang, H., Sievert, S., Liu, S., Charles, Z., Papailiopoulos, D., Wright, S.: Atomo: Communication-efficient learning via atomic sparsification. *Advances in Neural Information Processing Systems* **31** (2018)
38. Wang, J., Yuan, B., Rimanic, L., He, Y., Dao, T., Chen, B., Ré, C., Zhang, C.: Fine-tuning language models over slow networks using activation quantization with guarantees. *Advances in Neural Information Processing Systems* **35**, 19215–19230 (2022)