
Adaptive Gradient Quantization for Data-Parallel SGD

Fartash Faghri^{1,2*}

Iman Tabrizian^{1,2*}

Ilia Markov³

Dan Alistarh^{3,4}

Daniel M. Roy^{1,2}

Ali Ramezani-Kebrya²

¹University of Toronto ²Vector Institute ³IST Austria ⁴NeuralMagic

faghri@cs.toronto.edu iman.tabrizian@mail.utoronto.ca alir@vectorinstitute.ai

Abstract

Many communication-efficient variants of SGD use gradient quantization schemes. These schemes are often heuristic and fixed over the course of training. We empirically observe that the statistics of gradients of deep models change during the training. Motivated by this observation, we introduce two adaptive quantization schemes, ALQ and AMQ. In both schemes, processors update their compression schemes in parallel by efficiently computing sufficient statistics of a parametric distribution. We improve the validation accuracy by almost 2% on CIFAR-10 and 1% on ImageNet in challenging low-cost communication setups. Our adaptive methods are also significantly more robust to the choice of hyperparameters.

1 Introduction

Stochastic gradient descent (SGD) and its variants are currently the method of choice for training deep models. Yet, large datasets cannot always be trained on a single computational node due to memory and scalability limitations. Data-parallel SGD is a remarkably scalable variant, in particular on multi-GPU systems [1–10]. However, despite its many advantages, distribution introduces new challenges for optimization algorithms. In particular, data-parallel SGD has large communication cost due to the need to transmit potentially huge gradient vectors. Ideally, we want distributed optimization methods that match the performance of SGD on a single hypothetical super machine, while paying a negligible communication cost.

A common approach to reducing the communication cost in data-parallel SGD is gradient compression and quantization [4, 11–16]. In full-precision data-parallel SGD, each processor broadcasts its locally computed stochastic gradient vector at every iteration, whereas in quantized data-parallel SGD, each processor compresses its stochastic gradient before broadcasting. Current quantization methods are either designed heuristically or fixed prior to training. Convergence rates in a stochastic optimization problem are controlled by the trace of the gradient covariance matrix, which is referred




Figure 1: Changes in the average variance of normalized gradient coordinates in a ResNet-32 model trained on CIFAR-10. Colors distinguish different runs with different seeds. Learning rate is decayed by a factor of 10 twice at 40K and 60K iterations. The variance changes rapidly during the first epoch. The next noticeable change happens after the first learning rate drop and another one appears after the second drop.

*Equal contributions.

as the gradient variance in this paper [17]. As Fig. 1 shows, no fixed method can be optimal throughout the entire training because the distribution of gradients changes. A quantization method that is optimal at the first iteration will not be optimal after only a single epoch.

In this paper, we propose two adaptive methods for quantizing the gradients in data-parallel SGD. We study methods that are defined by a norm and a set of quantization levels. In Adaptive Level Quantization (ALQ), we minimize the excess variance of quantization given an estimate of the distribution of the gradients. In Adaptive Multiplier Quantization (AMQ), we minimize the same objective as ALQ by modelling quantization levels as exponentially spaced levels. AMQ solves for the optimal value of a single multiplier parametrizing the exponentially spaced levels.

1.1 Summary of contributions

- We propose two adaptive gradient quantization methods, ALQ and AMQ, in which processors update their compression methods in parallel.
- We establish an upper bound on the excess variance for any arbitrary sequence of quantization levels under general normalization that is tight in dimension, an upper bound on the expected number of communication bits per iteration, and strong convergence guarantees on a number of problems under standard assumptions. Our bounds hold for any adaptive method, including ALQ and AMQ.
- We improve the validation accuracy by almost 2% on CIFAR-10 and 1% on ImageNet in challenging low-cost communication setups. Our adaptive methods are significantly more robust to the choice of hyperparameters.²

1.2 Related work

Adaptive quantization has been used for speech communication and storage [18]. In machine learning, several biased and unbiased schemes have been proposed to compress networks and gradients. Recently, lattice-based quantization has been studied for distributed mean estimation and variance reduction [19]. In this work, we focus on unbiased and coordinate-wise schemes to compress gradients.

Alistarh et al. [20] proposed Quantized SGD (QSGD) focusing on the uniform quantization of stochastic gradients normalized to have unit Euclidean norm. Their experiments illustrate a similar quantization method, where gradients are normalized to have unit L^∞ norm, achieves better performance. We refer to this method as QSGDinf or Qinf in short. Wen et al. [15] proposed TernGrad, which can be viewed as a special case of QSGDinf with three quantization levels.

Ramezani-Kebrya et al. [21] proposed nonuniform quantization levels (NUQSGD) and demonstrated superior empirical results compared to QSGDinf. Horváth et al. [22] proposed natural compression and dithering schemes, where the latter is a special case of logarithmic quantization.

There have been prior attempts at adaptive quantization methods. Zhang et al. [23] proposed ZipML, which is an optimal quantization method if all points to be quantized are known a priori. To find the optimal sequence of quantization levels, a dynamic program is solved whose computational and memory cost is quadratic in the number of points to be quantized, which in the case of gradients would correspond to their dimension. For this reason, ZipML is impractical for quantizing on the fly, and is in fact used for (offline) dataset compression. They also proposed an approximation where a subsampled set of points is used and proposed to scan the data once to find the subset. However, as we show in this paper, this one-time scan is not enough as the distribution of stochastic gradients changes during the training.

Zhang et al. [24] proposed LQ-Net, where weights and activations are quantized such that the inner products can be computed efficiently with bitwise operations. Compared to LQ-Net, our methods do not need additional memory for encoding vectors. Concurrent with our work, Fu et al. [25] proposed to quantize activations and gradients by modelling them with Weibull distributions. In comparison, our proposed methods accommodate general distributions. Further, our approach does not require any assumptions on the upper bound of the gradients.

²Open source code: <http://github.com/tabrizian/learning-to-quantize>

Input: Local data, parameter vector (local copy) \mathbf{w}_t , learning rate α , and set of update steps \mathcal{U}

```

1 for  $t = 1$  to  $T$  do
2   if  $t \in \mathcal{U}$  then
3     for  $i = 1$  to  $M$  do
4       Compute sufficient statistics and update quantization levels  $\ell$ ;
5   for  $i = 1$  to  $M$  do
6     Compute  $g_i(\mathbf{w}_t)$ , encode  $c_{i,t} \leftarrow \text{ENCODE}_\ell(g_i(\mathbf{w}_t))$ , and broadcast  $c_{i,t}$ ;
7   for  $j = 1$  to  $M$  do
8     Receive  $c_{i,t}$  from each processor  $i$  and decode  $\hat{g}_i(\mathbf{w}_t) \leftarrow \text{DECODE}_\ell(c_{i,t})$ ;
9   Aggregate  $\mathbf{w}_{t+1} \leftarrow \mathbf{P}_\Omega(\mathbf{w}_t - \frac{\alpha}{M} \sum_{i=1}^M \hat{g}_i(\mathbf{w}_t))$ ;

```

Algorithm 1: Adaptive data-parallel SGD. Loops are executed in parallel on each machine. At certain steps, each processor computes sufficient statistics of a parametric distribution to estimate distribution of normalized coordinates.

2 Preliminaries: data-parallel SGD

Consider the problem of training a model parametrized by a high-dimensional vector $\mathbf{w} \in \mathbb{R}^d$. Let $\Omega \subseteq \mathbb{R}^d$ denote a closed and compact set. Our goal is to minimize $f : \Omega \rightarrow \mathbb{R}$. Assume we have access to unbiased stochastic gradients of f , which is g , such that $\mathbb{E}[g(\mathbf{w})] = \nabla f(\mathbf{w})$ for all $\mathbf{w} \in \Omega$.

The update rule for full-precision SGD is given by $\mathbf{w}_{t+1} = \mathbf{P}_\Omega(\mathbf{w}_t - \alpha g(\mathbf{w}_t))$ where \mathbf{w}_t is the current parameter vector, α is the learning rate, and \mathbf{P}_Ω is the Euclidean projection onto Ω . We consider data-parallel SGD, which is a synchronous and distributed framework consisting of M processors. Each processor receives gradients from all other processors and aggregates them. In data-parallel SGD with compression, gradients are compressed by each processor before transmission and decompressed before aggregation [20–23]. A stochastic compression method is unbiased if the vector after decompression is in expectation the same as the original vector.

3 Adaptive quantization

In this section, we introduce novel adaptive compression methods that adapt during the training (Algorithm 1). Let $\mathbf{v} \in \mathbb{R}^d$ be a vector we seek to quantize and $r_i = |v_i|/\|\mathbf{v}\|$ be its normalized coordinates for $i = 1, \dots, d$.³ Let $q_\ell(r) : [0, 1] \rightarrow [0, 1]$ denote a random quantization function applied to the normalized coordinate r using adaptable quantization levels, $\ell = [\ell_0, \dots, \ell_{s+1}]^\top$, where $0 = \ell_0 < \ell_1 < \dots < \ell_s < \ell_{s+1} = 1$. For $r \in [0, 1]$, let $\tau(r)$ denote the index of a level such that $\ell_{\tau(r)} \leq r < \ell_{\tau(r)+1}$. Let $\rho(r) = (r - \ell_{\tau(r)})/(\ell_{\tau(r)+1} - \ell_{\tau(r)})$ be the relative distance of r to level $\tau(r) + 1$. We define the random variable $h(r)$ such that $h(r) = \ell_{\tau(r)}$ with probability $1 - \rho(r)$ and $h(r) = \ell_{\tau(r)+1}$ with probability $\rho(r)$.

We define the quantization of \mathbf{v} as $Q_\ell(\mathbf{v}) \triangleq [q_\ell(v_1), \dots, q_\ell(v_d)]^\top$ where $q_\ell(v_i) = \|\mathbf{v}\| \cdot \text{sign}(v_i) \cdot h(r_i)$ and $\mathbf{h} = \{h(r_i)\}_{i=1, \dots, d}$ are independent random variables. The encoding, $\text{ENCODE}(\mathbf{v})$, of a stochastic gradient is the combined encoding of $\|\mathbf{v}\|$ using a standard floating point encoding along with an optimal encoding of $h(r_i)$ and binary encoding of $\text{sign}(v_i)$ for each coordinate i . The decoding, DECODE , recovers the norm, $h(r_i)$, and the sign. Additional details of the encoding method are described in Appendix D.




Figure 2: Random quantization of normalized gradient.

We define the variance of vector quantization to be the trace of the covariance matrix,

$$\mathbb{E}_{\mathbf{h}}[\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2] = \|\mathbf{v}\|^2 \sum_{i=1}^d \sigma^2(r_i), \quad (1)$$

where $\sigma^2(r) = \mathbb{E}[(q_\ell(r) - r)^2]$ is the variance of quantization for a single coordinate that is given by

$$\sigma^2(r) = (\ell_{\tau(r)+1} - r)(r - \ell_{\tau(r)}). \quad (2)$$

³In this section, we use $\|\cdot\|$ to denote a general L^q norm with $q \geq 1$ for simplicity.

Let \mathbf{v} be a random vector corresponding to a stochastic gradient and \mathbf{h} capture the randomness of quantization for this random vector as defined above. We define two minimization problems, expected variance and expected normalized variance minimization:

$$\min_{\ell \in \mathcal{L}} \mathbb{E}_{\mathbf{v}, \mathbf{h}} [\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2] \quad \text{and} \quad \min_{\ell \in \mathcal{L}} \mathbb{E}_{\mathbf{v}, \mathbf{h}} [\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2 / \|\mathbf{v}\|^2],$$

where $\mathcal{L} = \{\ell : \ell_j \leq \ell_{j+1}, \forall j, \ell_0 = 0, \ell_{s+1} = 1\}$ denotes the set of feasible solutions. We first focus on the problem of minimizing the expected normalized variance and then extend our methods to minimize the expected variance in Section 3.4. Let $F(r)$ denote the marginal cumulative distribution function (CDF) of a normalized coordinate r . Assuming normalized coordinates r_i are i.i.d. given $\|\mathbf{v}\|$, the expected normalized variance minimization can be written as

$$\min_{\ell \in \mathcal{L}} \Psi(\ell), \text{ where } \Psi(\ell) \triangleq \sum_{j=0}^s \int_{\ell_j}^{\ell_{j+1}} \sigma^2(r) dF(r). \quad (3)$$

The following theorem suggests that solving (3) is challenging in general; however, the sub-problem of optimizing a single level given other levels can be solved efficiently in closed form. Proofs are provided in Appendix B.

Theorem 1 (Expected normalized variance minimization). *Problem (3) is nonconvex in general. However, the optimal solution to minimize one level given other levels, $\min_{\ell_i} \Psi(\ell)$, is given by $\ell_i^* = \beta(\ell_{i-1}, \ell_{i+1})$, where*

$$\beta(a, c) = F^{-1} \left(F(c) - \int_a^c \frac{r-a}{c-a} dF(r) \right). \quad (4)$$

3.1 ALQ: Adapting individual levels using coordinate descent

Using the single level update rule in Eq. (4) we iteratively adapt individual levels to minimize the expected normalized variance in (3). We denote quantization levels at iteration t by $\ell(t)$ starting from $t = 0$. The update rule is

$$\ell_j(t+1) = \beta(\ell_{j-1}(t), \ell_{j+1}(t)) \quad \forall j = 1, \dots, s. \quad (5)$$

Performing the update rule above sequentially over coordinates j is a form of coordinate descent (CD) that is guaranteed to converge to a local minima. CD is particularly interesting because it does not involve any projection step to the feasible set \mathcal{L} . In practice, we initialize the levels with either uniform levels [20] or exponentially spaced levels proposed in [21]. We observe that starting from either initialization CD converges in small number of steps (less than 10).

3.2 Gradient descent

Computing $\nabla \Psi$ using Leibniz's rule [26], the gradient descent (GD) algorithm to solve (3) is based on the following update rule:

$$\begin{aligned} \ell_j(t+1) &= \mathbf{P}_{\mathcal{L}} \left(\ell_j(t) - \eta(t) \frac{\partial \Psi(\ell(t))}{\partial \ell_j} \right) \\ \frac{\partial \Psi(\ell(t))}{\partial \ell_j} &= \int_{\ell_{j-1}(t)}^{\ell_j(t)} (r - \ell_{j-1}(t)) dF(r) - \int_{\ell_j(t)}^{\ell_{j+1}(t)} (\ell_{j+1}(t) - r) dF(r) \end{aligned} \quad (6)$$

for $t = 0, 1, \dots$ and $j = 1, \dots, s$. Note that the projection step in Eq. (6) is itself a convex optimization problem. We propose a projection-free modification of GD update rule to systematically ensure $\ell \in \mathcal{L}$. Let $\delta_j(t) = \min\{\ell_j(t) - \ell_{j-1}(t), \ell_{j+1}(t) - \ell_j(t)\}$ denote the minimum distance between two neighbouring levels at iteration t for $j = 1, \dots, s$. If the change in level j is bounded by $\delta_j(t)/2$, it is guaranteed that $\ell \in \mathcal{L}$. We propose to replace Eq. (6) with the following update rule:

$$\ell_j(t+1) = \ell_j(t) - \text{sign} \left(\frac{\partial \Psi(\ell(t))}{\partial \ell_j} \right) \min \left\{ \eta(t) \left| \frac{\partial \Psi(\ell(t))}{\partial \ell_j} \right|, \frac{\delta_j(t)}{2} \right\}. \quad (7)$$

3.3 AMQ: Exponentially spaced levels

We now focus on $\ell = [-1, -p, \dots, -p^s, p^s, \dots, p, 1]^\top$, *i.e.*, exponentially spaced levels with symmetry. We can update p efficiently by gradient descent using the first order derivative

$$\frac{1}{2} \frac{d\Psi(p)}{dp} = \int_0^{p^s} 2sp^{2s-1} dF(r) + \sum_{j=0}^{s-1} \int_{p^{j+1}}^{p^j} ((jp^{j-1} + (j+1)p^j)r - (2j+1)p^{2j}) dF(r). \quad (8)$$

3.4 Expected variance minimization

In this section, we consider the problem of minimizing the expected variance of quantization:

$$\min_{\ell \in \mathcal{L}} \mathbb{E}_{\mathbf{v}, \mathbf{h}} [\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2]. \quad (9)$$

To solve the expected variance minimization problem, suppose that we observe N stochastic gradients $\{\mathbf{v}_1, \dots, \mathbf{v}_N\}$. Let $F_n(r)$ and $p_n(r)$ denote the CDF and PDF of normalized coordinate conditioned on observing $\|\mathbf{v}_n\|$, respectively. By taking into account randomness in $\|\mathbf{v}\|$ and using the law of total expectation, an approximation of the expected variance in (9) is given by

$$\mathbb{E}[\|Q_s(\mathbf{v}) - \mathbf{v}\|_2^2] \approx \frac{1}{N} \sum_{n=1}^N \|\mathbf{v}_n\|^2 \sum_{j=0}^s \int_{\ell_j}^{\ell_{j+1}} \sigma^2(r) dF_n(r). \quad (10)$$

The optimal levels to minimize Eq. (10) are a solution to the following problem:

$$\ell^* = \arg \min_{\ell \in \mathcal{L}} \sum_{n=1}^N \|\mathbf{v}_n\|^2 \sum_{j=0}^s \int_{\ell_j}^{\ell_{j+1}} \sigma^2(r) dF_n(r) = \arg \min_{\ell \in \mathcal{L}} \sum_{j=0}^s \int_{\ell_j}^{\ell_{j+1}} \sigma^2(r) d\bar{F}(r),$$

where $\ell^* = [\ell_1^*, \dots, \ell_s^*]^\top$ and $\bar{F}(r) = \sum_{n=1}^N \gamma_n F_n(r)$ is the weighted sum of the conditional CDFs with $\gamma_n = \|\mathbf{v}_n\|^2 / \sum_{n=1}^N \|\mathbf{v}_n\|^2$. Note that we can accommodate both normal and truncated normal distributions by substituting associated expressions into $p_n(r)$ and $F_n(r)$. Exact update rules and analysis of computational complexity of ALQ, GD, and AMQ are discussed in Appendix C.

4 Theoretical guarantees

One can alternatively design quantization levels to minimize the worst-case variance. However, compared to an optimal scheme, this worst-case scheme increases the expected variance by $\Omega(d)$, which is prohibitive in deep networks. We quantify the gap in Appendix E. Proofs are in appendices.

A stochastic gradient has a *second-moment upper bound* B when $\mathbb{E}[\|g(\mathbf{w})\|_2^2] \leq B$ for all $\mathbf{w} \in \Omega$. Similarly, it has a *variance upper bound* σ^2 when $\mathbb{E}[\|g(\mathbf{w}) - \nabla f(\mathbf{w})\|_2^2] \leq \sigma^2$ for all $\mathbf{w} \in \Omega$.

We consider a general adaptively quantized SGD (AQSGD) algorithm, described in Algorithm 1, where compression schemes are updated over the course of training.⁴ Many convergence results in stochastic optimization rely on a variance bound. We establish such a variance bound for our adaptive methods. Further, we verify that these optimization results can be made to rely only on the average variance. In the following, we provide theoretical guarantees for AQSGD algorithm, obtain variance and code-length bounds, and convergence guarantees for convex, nonconvex, and momentum-based variants of AQSGD.

The analysis of nonadaptive methods in [20–23] can be considered as special cases of our theorems with fixed levels over the course of training. A naive adoption of available convergence guarantees results in having worst-case variance bounds over the course of training. In this paper, we show that an average variance bound can be applied on a number of problems. Under general normalization, we first obtain variance upper bound for arbitrary levels, in particular, for those obtained adaptively.

Theorem 2 (Variance bound). *Let $\mathbf{v} \in \mathbb{R}^d$ and $q \geq 1$. The quantization of \mathbf{v} under L^q normalization satisfies $\mathbb{E}[Q_\ell(\mathbf{v})] = \mathbf{v}$. Furthermore, we have*

$$\mathbb{E}[\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2] \leq \epsilon_Q \|\mathbf{v}\|_2^2, \quad (11)$$

⁴Our results hold for any adaptive method, including ALQ and AMQ.

where $\epsilon_Q = \frac{(\ell_{j^*+1}/\ell_{j^*}-1)^2}{4(\ell_{j^*+1}/\ell_{j^*})} + \inf_{0 < p < 1} K_p \ell_1^{(2-p)} d^{\frac{2-p}{\min\{q, 2\}}}$ with $j^* = \arg \max_{1 \leq j \leq s} \ell_{j+1}/\ell_j$ and $K_p = \left(\frac{1}{2-p}\right) \left(\frac{1-p}{2-p}\right)^{(1-p)}$.

Theorem 2 implies that if $g(\mathbf{w})$ is a stochastic gradient with a second-moment bound η , then $Q_\ell(g(\mathbf{w}))$ is a stochastic gradient with a variance upper bound $\epsilon_Q \eta$. Note that, as long as the maximum ratio of two consecutive levels does not change, the variance upper bound decreases with the number of quantization levels. In addition, our bound matches the known $\Omega(\sqrt{d})$ lower bound in [27].

Theorem 3 (Code-length bound). *Let $\mathbf{v} \in \mathbb{R}^d$ and $q \geq 1$. The expectation $\mathbb{E}[|\text{ENCODE}(\mathbf{v})|]$ of the number of communication bits needed to transmit $Q_\ell(\mathbf{v})$ under L^q normalization is bounded by*

$$\mathbb{E}[|\text{ENCODE}(\mathbf{v})|] \leq b + n_{\ell_1, d} + d(H(L) + 1) \leq b + n_{\ell_1, d} + d(\log_2(s + 2) + 1), \quad (12)$$

where b is a constant, $n_{\ell_1, d} = \min\{\ell_1^{-q} + \frac{d^{1-1/q}}{\ell_1}, d\}$, $H(L)$ is the entropy of L in bits, and L is a random variable with the probability mass function given by

$$\Pr(\ell_j) = \int_{\ell_{j-1}}^{\ell_j} \frac{r - \ell_{j-1}}{\ell_j - \ell_{j-1}} dF(r) + \int_{\ell_j}^{\ell_{j+1}} \frac{\ell_{j+1} - r}{\ell_{j+1} - \ell_j} dF(r)$$

for $j = 1, \dots, s$. In addition, we have

$$\Pr(\ell_0 = 0) = \int_0^{\ell_1} \frac{1 - r}{\ell_1} dF(r) \text{ and } \Pr(\ell_{s+1} = 1) = \int_{\ell_s}^1 \frac{r - \ell_s}{1 - \ell_s} dF(r).$$

Theorem 3 provides a bound on the expected number of communication bits to encode the quantized stochastic gradients. As expected, the upper bound in (12) increases monotonically with d and s .

We can combine variance and code-length upper bounds and obtain convergence guarantees for AQSGD when applied to various learning problems where we have convergence guarantees for full-precision SGD under standard assumptions.

Let $\{\ell_1, \dots, \ell_K\}$ denote the set of quantization levels that AQSGD experiences on the optimization trajectory. Suppose that ℓ_k is used for T_k iterations with $\sum_{k=1}^K T_k = T$. For each particular ℓ_k , we can obtain corresponding variance bound $\epsilon_{Q,k}$ by substituting ℓ_k into (11). Then the average variance upper bound is given by $\overline{\epsilon_Q} = \sum_{k=1}^K T_k \epsilon_{Q,k} / T$. For each particular ℓ_k , we can obtain corresponding expected code-length bound $N_{Q,k}$ by substituting random variable L_k into (12). The average expected code-length bound is given by $\overline{N_Q} = \sum_{k=1}^K T_k N_{Q,k} / T$.

On convex problems, convergence guarantees can be established along the lines of [17, Theorems 6.1].

Theorem 4 (AQSGD for nonsmooth convex optimization). *Let $f : \Omega \rightarrow \mathbb{R}$ denote a convex function and let $R^2 \triangleq \sup_{\mathbf{w} \in \Omega} \|\mathbf{w} - \mathbf{w}_0\|_2^2$. Let $\hat{B} = (1 + \overline{\epsilon_Q})B$ and $f^* = \inf_{\mathbf{w} \in \Omega} f(\mathbf{w})$. Suppose that AQSGD is executed for T iterations with a learning rate $\alpha = RM/(\hat{B}\sqrt{T})$ on M processors, each with access to independent stochastic gradients of f with a second-moment bound B , such that quantization levels are updated K times where ℓ_k with variance bound $\epsilon_{Q,k}$ and code-length bound $N_{Q,k}$ is used for T_k iterations. Then AQSGD satisfies $\mathbb{E} \left[f \left(\frac{1}{T} \sum_{t=0}^T \mathbf{w}_t \right) \right] - f^* \leq R\hat{B}/(M\sqrt{T})$.*

In addition, AQSGD requires at most $\overline{N_Q}$ communication bits per iteration in expectation.

In Appendix H and Appendix I, we obtain convergence guarantees on nonconvex problems and for momentum-based variants of AQSGD under standard assumptions, respectively. Theoretical guarantees for levels with symmetry are established in Appendix J.

5 Experimental evaluation

In this section, we showcase the effectiveness of our adaptive quantization methods in speeding up training deep models. We compare our methods to the following baselines: single-GPU SGD (SGD), full-precision multi-GPU SGD (SuperSGD), uniform levels under L^∞ normalization (QSGDinf) [20], ternary levels under L^∞ normalization (TRN) [15], and exponential levels under L^2 normalization with exponential factor $p = 0.5$ (NUQSGD) [21, 22]. We present results for the following variations of

Table 1: Validation accuracy on CIFAR-10 and ImageNet using 3 bits (except for SuperSGD and TRN) with 4 GPUs.

Quantization Method	ResNet-110 on CIFAR-10	ResNet-32 on CIFAR-10	ResNet-18 on ImageNet
Bucket Size	16384	8192	8192
SuperSGD	93.86% \pm 0.08	92.26% \pm 0.04	68.93% \pm 0.05
NUQSGD [21, 22]	84.60% \pm 0.04	83.73% \pm 0.08	33.36% \pm 0.07
QSGDinf [20]	91.52% \pm 0.07	89.95% \pm 0.02	66.35% \pm 0.04
TRN [15]	90.72% \pm 0.06	89.65% \pm 0.05	62.76% \pm 0.06
ALQ	93.24% \pm 0.06	91.30% \pm 0.07	67.72% \pm 0.07
ALQ-N	93.14% \pm 0.05	91.96% \pm 0.04	65.64% \pm 0.07
AMQ	92.82% \pm 0.04	91.10% \pm 0.05	64.82% \pm 0.05
AMQ-N	92.88% \pm 0.02	91.03% \pm 0.08	66.75% \pm 0.05




Figure 3: Validation loss on CIFAR-10 and ImageNet. All methods use 3 bits except for SuperSGD and TRN. Bucket size for ResNet-110 trained on CIFAR-10 is 16384, for ResNet-32 is 8192, and for ResNet-18 on ImageNet is 8192.

our proposed methods: ALQ and AMQ (with norm adjustments in Section 3.4), and their normalized variations ALQ-N and AMQ-N (Sections 3.1 and 3.3). We present full training results on ImageNet in Appendix K along with additional experimental details.

We compare methods in terms of the number of training iterations that is independent of a particular distributed setup. In Table 1, we present results for training ResNet-32 and ResNet-110 [28] on CIFAR-10 [29], and ResNet-18 on ImageNet [30]. We simulate training with 4-GPUs on a single GPU by quantizing and dequantizing the gradient from 4 mini-batches in each training iteration. These simulations allow us to compare the performance of quantization methods to the hypothetical full-precision SuperSGD.

All quantization methods studied in this section share two hyper-parameters: the number of bits (\log_2 of number of quantization levels) and a bucket size. A common trick used in normalized quantization is to encode and decode a high-dimensional vector in buckets such that each coordinate is normalized by the norm of its corresponding bucket instead of the norm of the entire vector [20]. The bucket size controls the tradeoff between extra communication cost and loss of precision. With a small bucket size, there are more bucket norms to be communicated, while with a large bucket size, we lose numerical precision as a result of dividing each coordinate by a large number. In Section 5.1, we provide an empirical study of the hyperparameters.

Matching the accuracy of SuperSGD. Using only 3 bits (8 levels), our adaptive methods match the performance of SuperSGD on CIFAR-10 and close the gap on ImageNet (bold in Table 1). Our most flexible method, ALQ, achieves the best overall performance on ImageNet and the gap on CIFAR-10 with ALQ-N is less than 0.3%. There is at least 1.4% gap between our best performing method and previous work in training each model. To the best of our knowledge, matching the validation loss of SuperSGD has not been achieved in any previous work using only 3 bits. Fig. 3 shows the test loss and Fig. 4 shows the average gradient variance where the average is taken over gradient coordinates. Our adaptive methods successfully achieve lower variance during training.

Comparison on the trajectory of SGD. Fig. 5 shows the average variance on the optimization trajectory of single-GPU without quantization. This graph provides a more fair comparison of the




Figure 4: Variance on CIFAR-10 and ImageNet. All methods use 3 bits except for SuperSGD and TRN. Bucket size for ResNet-110 trained on CIFAR-10 is 16384, for ResNet-32 is 8192, and for ResNet-18 on ImageNet is 8192.




Figure 5: Variance (no train) on CIFAR-10 and ImageNet. All methods use 3 bits except for SuperSGD and TRN. Bucket size for ResNet-110 trained on CIFAR-10 is 16384, for ResNet-32 is 8192, and for ResNet-18 on ImageNet is 8192.

quantization error of different methods decoupled from their impact on the optimization trajectory. ALQ effectively finds an improved set of levels that reduce the variance in quantization. ALQ matches the variance of SuperSGD on Resnet-110 (Fig. 5b). In Figs. 5b and 5c, the variance of QSGDinf is as high as TRN in the first half of training. This shows that extra levels (8 uniform levels) do not perform better unless designed carefully. As expected, the variance of SuperSGD is always smaller than the variance of SGD by a constant factor of the number of GPUs.

Negligible computational overhead. Our adaptive methods have similar per-step computation and communication cost compared to previous methods. On ImageNet, we save at least 60 hours from 95 hours of training and add only an additional cost of at most 10 minutes in total to adapt quantization. For bucket sizes 8192 and 16384 and 3–8 bits used in our experiments, the per-step cost relative to SuperSGD (32-bits) is 21–25% for ResNet-18 on ImageNet and 32–36% for ResNet-50. That is the same as the cost of NUQSGD and QSGDinf without additional coding or pruning with the same number of bits and bucket sizes. The cost of the additional update specific to ALQ is 0.4–0.5% of the total training time. In Appendix K.3, we provide tables with detailed timing results for varying bucket sizes and bits.

5.1 Hyperparameter studies

Fig. 6 shows quantization levels for each method at the end of training ResNet-32 on CIFAR-10. The quantization levels for our adaptive methods are more concentrated near zero. In Figs. 7a and 7b, we study the impact of the bucket size and number of bits on the best validation accuracy achieved by quantization methods.

Adaptive levels are the best quantization methods across all values of bucket size and number of bits. ALQ and ALQ-N are the best performing methods across all values of bucket size and number of bits. The good performance of ALQ-N is unexpected as it suggests quantization for vectors with different norms can be shared. In practice, ALQ-N is easier to implement and faster to update compared to ALQ. We observe a similar relation between AMQ and AMQ-N methods. Adaptive multiplier methods show inferior performance to adaptive level methods as the bucket size significantly grows (above 10^4) or shrinks (below 100) as well as for very few bits (2). Note that there exists a known generalization gap between SGD and SuperSGD in ResNet-110 that can be closed by extensive hyperparameter tuning [31]. Our adaptive methods reduce this gap with standard hyperparameters.




Figure 6: Quantization levels at the end of training ResNet-32 on CIFAR-10.




Figure 7: Effect of bucket size and number of bits on validation accuracy when training ResNet-8 on CIFAR-10

Table 2: Validation accuracy of ResNet32 on CIFAR-10 using 3 quantization bits (except for SuperSGD and TRN) and bucket size 16384.

Method	16 GPUs	32 GPUs
SuperSGD	92.17% \pm 0.08	92.19% \pm 0.04
NUQSGD	85.82% \pm 0.03	86.36% \pm 0.01
QSGDinf	89.61% \pm 0.03	89.81% \pm 0.05
TRN	88.68% \pm 0.10	90.22% \pm 0.05
ALQ	91.91% \pm 0.06	91.89% \pm 0.07
ALQ-N	92.07% \pm 0.04	91.83% \pm 0.03
AMQ	91.58% \pm 0.05	91.38% \pm 0.06
AMQ-N	91.41% \pm 0.08	91.40% \pm 0.02

Bucket size significantly impacts non-adaptive methods. For bucket size 100 and 3 bits, NUQSGD performs nearly as good as adaptive methods but quickly loses accuracy as the bucket size grows or shrinks. QSGDinf stays competitive for a wider range of bucket sizes but still loses accuracy faster than other methods. This shows the impact of bucketing as an understudied trick in evaluating quantization methods.

Adaptive methods successfully scale to large number of GPUs. Table 2 shows the result of training CIFAR-10 on ResNet-32 using 16 and 32 GPUs. Note that with 32 GPUs, TRN is achieving almost the accuracy of SuperSGD with only 3 quantization levels, which is expected because TRN is unbiased and the variance of aggregated gradients decreases linearly with the number of GPUs.

6 Conclusions

To reduce communication costs of data-parallel SGD, we introduce two adaptively quantized methods, ALQ and AMQ, to learn and adapt gradient quantization method on the fly. In addition to quantization method, in both methods, processors learn and adapt their coding methods in parallel by efficiently computing sufficient statistics of a parametric distribution. We establish tight upper bounds on the excessive variance for any arbitrary sequence of quantization levels under general normalization and on the expected number of communication bits per iteration. Under standard assumptions, we establish a number of convergence guarantees for our adaptive methods. We demonstrate the superiority of ALQ and AMQ over nonadaptive methods empirically on deep models and large datasets.

Broader impact

This work provides additional understanding of statistical behaviour of deep machine learning models. We aim to train deep models using popular SGD algorithm as fast as possible without compromising learning outcome. As the amount of data gathered through web and a plethora of sensors deployed everywhere (e.g., IoT applications) is drastically increasing, the design of efficient machine learning algorithms that are capable of processing large-scale data in a reasonable time can improve everyone’s quality of life. Our compression schemes can be used in Federated Learning settings, where a deep

model is trained on data distributed among multiple owners without exposing that data. Developing privacy-preserving learning algorithms is an integral part of responsible and ethical AI. However, the long-term impacts of our schemes may depend on how machine learning is used in society.

Acknowledgement

The authors would like to thank Blair Bilodeau, David Fleet, Mufan Li, and Jeffrey Negrea for helpful discussions. FF was supported by OGS Scholarship. DA and IM were supported the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 805223 ScaleML). DMR was supported by an NSERC Discovery Grant. ARK was supported by NSERC Postdoctoral Fellowship. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute.⁵

References

- [1] M. Zinkevich, M. Weimer, L. Li, and A. J. Smola. Parallelized stochastic gradient descent. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [2] R. Bekkerman, M. Bilenko, and J. Langford. *Scaling up machine learning: Parallel and distributed approaches*. Cambridge University Press, 2011.
- [3] B. Recht, C. Ré, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2011.
- [4] J. Dean, G. Corrado, R. Monga K. Chen, M. Devin, M. Mao, M. Ranzato, A. Senior, P. Tucker, K. Yang, Q. V. Le, and A. Y. Ng. Large scale distributed deep networks. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [5] A. Coates, B. Huval, T. Wang, D. Wu, B. Catanzaro, and A. Ng. Deep learning with COTS HPC systems. In *Proc. International Conference on Machine Learning (ICML)*, 2013.
- [6] T. Chilimbi, Y. Suzue J. Apacible, and K. Kalyanaraman. Project adam: Building an efficient and scalable deep learning training system. In *Proc. USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2014.
- [7] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, A. Ahmed, V. Josifovski, J. Long, E. J. Shekita, and B.-Y. Su. Scaling distributed machine learning with the parameter server. In *Proc. USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2014.
- [8] J. C. Duchi, S. Chaturapruek, and C. Ré. Asynchronous stochastic convex optimization. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [9] E. P. Xing, Q. Ho, W. Dai, J. K. Kim, J. Wei, S. Lee, X. Zheng, P. Xie, A. Kumar, and Y. Y. Petuum. Petuum: A new platform for distributed machine learning on big data. *IEEE transactions on Big Data*, 1(2):49–67, 2015.
- [10] S. Zhang, A. E. Choromanska, and Y. LeCun. Deep learning with elastic averaging SGD. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [11] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs. In *Proc. INTERSPEECH*, 2014.
- [12] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan. Deep learning with limited numerical precision. In *Proc. International Conference on Machine Learning (ICML)*, 2015.
- [13] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, and M. Devin. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv:1603.04467, 2016.

⁵www.vectorinstitute.ai/#partners

- [14] S. Zhou, Y. Wu, Z. Ni, X. Zhou, H. Wen, and Y. Zou. DoReFa-Net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv:1606.06160*, 2016.
- [15] W. Wen, C. Xu, F. Yan, C. Wu, Y. Wang, Y. Chen, and H. Li. TernGrad: Ternary gradients to reduce communication in distributed deep learning. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [16] J. Bernstein, Y.-X. Wang, K. Azizzadenesheli, and A. Anandkumar. signSGD: Compressed optimisation for non-convex problems. In *Proc. International Conference on Machine Learning (ICML)*, 2018.
- [17] S. Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–358, 2015.
- [18] P. Cummiskey, N. S. Jayant, and J. L. Flanagan. Adaptive quantization in differential PCM coding of speech. *Bell System Technical Journal*, 52(7):1105–1118, 1973.
- [19] D. Alistarh, S. Ashkboos, and P. Davies. Distributed mean estimation with optimal error bounds. *arXiv:2002.09268v2*, 2020.
- [20] D. Alistarh, D. Grubic, J. Z. Li, R. Tomioka, and M. Vojnovic. QSGD: Communication-efficient SGD via gradient quantization and encoding. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [21] A. Ramezani-Kebrya, F. Faghri, and D. M. Roy. NUQSGD: Improved communication efficiency for data-parallel SGD via nonuniform quantization. *arXiv preprint arXiv:1908.06077v1*, 2019.
- [22] S. Horváth, C.-Y. Ho, L. Horváth, A. N. Sahu, M. Canini, and P. Richtárik. Natural compression for distributed deep learning. *arXiv:1905.10988v1*, 2019.
- [23] H. Zhang, J. Li, K. Kara, D. Alistarh, J. Liu, and C. Zhang. ZipML: Training linear models with end-to-end low precision, and a little bit of deep learning. In *Proc. International Conference on Machine Learning (ICML)*, 2017.
- [24] D. Zhang, J. Yang, D. Ye, and G. Hua. LQ-Nets: Learned quantization for highly accurate and compact deep neural networks. In *Proc. European Conference on Computer Vision (ECCV)*, 2018.
- [25] F. Fu, Y. Hu, Y. He, J. Jiang, Y. Shao, C. Zhang, and B. Cui. Don’t waste your bits! squeeze activations and gradients for deep neural networks via TINYSRIPT. In *Proc. International Conference on Machine Learning (ICML)*, 2020.
- [26] M. H. Protter and C. B. Morrey. *Intermediate Calculus*. Springer, 1985.
- [27] A. Ramezani-Kebrya, F. Faghri, I. Markov, V. Aksenov, D. Alistarh, and D. M. Roy. NUQSGD: Provably communication-efficient data-parallel SGD via nonuniform quantization. *Technical Report*.
- [28] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [29] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [30] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [31] Dami Choi, Christopher J. Shallue, Zachary Nado, Jaehoon Lee, Chris J. Maddison, and George E. Dahl. On Empirical Comparisons of Optimizers for Deep Learning. *arXiv e-prints*, art. arXiv:1910.05446, October 2019.
- [32] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. WILEY, 2006.

- [33] S. Ghadimi and G. Lan. Stochastic first- and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- [34] T. Yang, Q. Lin, and Z. Li. Unified convergence analysis of stochastic momentum methods for convex and non-convex optimization. arXiv:1604.03257v2, 2016.
- [35] B. T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1–17, 1964.
- [36] Y. Nesterov. A method of solving a convex programming problem with convergence $O(1/k^2)$. *Soviet Mathematics Doklady*, 27(2):372–376, 1983.

A CDF and its inverse

A.1 Normal distribution

The probability density function (PDF) for $X \sim \mathcal{N}(\mu, \sigma^2)$ is defined as

$$p_{\mathcal{N}}(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad (13)$$

and the cumulative distribution function (CDF) defined as

$$F_{\mathcal{N}}(x) = \Phi\left(\frac{x-\mu}{\sigma}\right) = 1 - Q\left(\frac{x-\mu}{\sigma}\right) = \frac{1}{2}\left(1 + \operatorname{erf}\left(\frac{x-\mu}{\sqrt{2}\sigma}\right)\right), \quad (14)$$

where

$$\begin{aligned} \Phi(x) &= \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} \exp(-u^2/2) du, \\ Q(x) &= \int_x^{\infty} \frac{1}{\sqrt{2\pi}} \exp(-u^2/2) du, \\ \operatorname{erf}(x) &= 2 \int_0^x \frac{1}{\sqrt{\pi}} \exp(-u^2) du. \end{aligned}$$

The inverse of CDF for the normal distribution is given by

$$F_{\mathcal{N}}^{-1}(y) = \sigma\Phi^{-1}(y) + \mu = \sqrt{2}\sigma\operatorname{erf}^{-1}(2y-1) + \mu. \quad (15)$$

Various approximations of Eq. (14) and Eq. (15) are available in the literature.

A.2 Truncated normal distribution

The probability density function (PDF) of a truncated normal distribution that lies within the interval (a, b) with $-\infty < a < b < \infty$ is defined as

$$p_{\mathcal{T}}(x; a, b) = \frac{p_{\mathcal{N}}(x)}{\sigma(F_{\mathcal{N}}(b) - F_{\mathcal{N}}(a))}, \quad (16)$$

where $p_{\mathcal{N}}(\cdot)$ is defined in Eq. (13) and the cumulative distribution function (CDF) is defined as

$$\begin{aligned} F_{\mathcal{T}}(x; a, b) &= \frac{F_{\mathcal{N}}(x) - F_{\mathcal{N}}(a)}{F_{\mathcal{N}}(b) - F_{\mathcal{N}}(a)} \\ &= \frac{\Phi((x-\mu)/\sigma) - \Phi((a-\mu)/\sigma)}{\Phi((b-\mu)/\sigma) - \Phi((a-\mu)/\sigma)}, \end{aligned} \quad (17)$$

where $F_{\mathcal{N}}(\cdot)$ and $\Phi(\cdot)$ are defined in Eq. (14). Note that the *mean* and *variance* of a random variable with truncated normal distribution are not μ and σ^2 based on our notation. The mean and variance depend on the interval (a, b) , which is clear in contexts that we use.

The inverse of CDF for truncated normal distribution is given by

$$F_{\mathcal{T}}^{-1}(y; a, b) = F_{\mathcal{N}}^{-1}(\bar{y}) = \sigma\Phi^{-1}(\bar{y}) + \mu, \quad (18)$$

where $\bar{y} = (F_{\mathcal{N}}(b) - F_{\mathcal{N}}(a))y + F_{\mathcal{N}}(a)$ and $F_{\mathcal{N}}^{-1}(\cdot)$, $\Phi^{-1}(\cdot)$ are defined in Eq. (15).

B Expected normalized variance minimization

B.1 Theorem 1

We prove Theorem 1 in two steps in Proposition 1 and Proposition 2.

Let R denote a random variable with probability density function (PDF) p and cumulative distribution function (CDF) F . To show that problem (3) is nonconvex, we first focus on the problem of optimizing two levels $\min_{a,b} Q(a,b)$ where

$$Q(a,b) = \int_0^a (a-r)r \, dF(r) + \int_a^b (b-r)(r-a) \, dF(r) + \int_b^1 (1-r)(r-b) \, dF(r) \quad (19)$$

in the range $0 \leq a \leq b \leq 1$.

Proposition 1. *The function $Q(a,b)$ is nonconvex in general. It becomes convex if for all $0 \leq a \leq b \leq 1$, we have*

$$b(1-a)p(a)p(b) > (F(b) - F(a))^2. \quad (20)$$

Proof. Using Leibniz's rule [26], we have

$$\nabla^2 Q = \begin{bmatrix} bp(a) & F(a) - F(b) \\ F(a) - F(b) & (1-a)p(b) \end{bmatrix}. \quad (21)$$

We can find the eigenvalues of $\nabla^2 Q$ by solving $|\nabla^2 Q - \lambda \mathbf{I}| = 0$, which leads to the following quadratic equation:

$$\lambda^2 - (bp(a) + (1-a)p(b))\lambda + (b(1-a)p(a)p(b) - (F(b) - F(a))^2) = 0. \quad (22)$$

We note that (20) is sufficient to guarantee $\nabla^2 Q \succeq \mathbf{0}$. \square

Corollary 1. *The sufficient condition (20) is satisfied if R is uniformly distributed in the range $[0, 1]$.*

We now solve the problem of optimizing a single level, i.e., $\min_b Q(b)$ where

$$Q(b) = \int_a^b (b-r)(r-a) \, dF(r) + \int_b^c (c-r)(r-b) \, dF(r). \quad (23)$$

Proposition 2. *The optimal solution to minimize $Q(b)$ is given by⁶*

$$b^* = F^{-1} \left(F(c) - \int_a^c \frac{r-a}{c-a} \, dF(r) \right). \quad (24)$$

Proof. Using Leibniz's rule, we have

$$\begin{aligned} \frac{dQ}{db} &= \int_a^b (r-a) \, dF(r) - \int_b^c (c-r) \, dF(r), \\ \frac{d^2Q}{db^2} &= (c-a)p(b). \end{aligned}$$

We note that Q is convex so we can find the closed-form optimal solution through satisfying the first order optimality condition. \square

Corollary 2. *In the special case with $a = 0$ and $c = 1$, the optimal solution to minimize $Q(b)$ is given by*

$$b^* = F^{-1}(1 - \mathbb{E}[R]).$$

For the special case of a truncated normal, the inner integral is evaluated as

$$\int_a^c \frac{r-a}{c-a} \, dF_{\mathcal{T}}(r) = \frac{\mu-a}{c-a} (F_{\mathcal{T}}(c) - F_{\mathcal{T}}(a)) - \frac{\sigma^2}{c-a} (p_{\mathcal{T}}(c) - p_{\mathcal{T}}(a)),$$

where $F_{\mathcal{T}}$ and $p_{\mathcal{T}}$, the CDF and PDF of the truncated normal, are defined in Appendix A.

⁶If F is not a one-to-one function, b^* can be any solution that satisfies $(c-a) \int_b^c \, dF(r) = \int_a^c (r-a) \, dF(r)$.

B.2 Projected Gradient Descent

For the special case of a normal or truncated normal distribution, the gradient of the expected normalized variance used in Section 3.2 is:

$$\begin{aligned} \frac{\partial \Psi(\ell(t))}{\partial \ell_j} &= (\mu - \ell_{j-1}(t)) (F(\ell_j(t)) - F(\ell_{j-1}(t))) + \sigma^2 (p(\ell_{j-1}(t)) - p(\ell_{j+1}(t))) \\ &\quad + (\mu - \ell_{j+1}(t)) (F(\ell_{j+1}(t)) - F(\ell_j(t))). \end{aligned} \quad (25)$$

B.3 Symmetric Levels

In this section, we introduce quantization method with symmetrical levels. This is particularly useful when the estimated PDF of normalized coordinates is an even function, which is the case for normal distribution with zero mean. Let $(-\ell_{s+1}, \dots, -\ell_1, \ell_1, \dots, \ell_{s+1})$ denote a sequence of symmetrical quantization levels w.r.t. 0 where $0 < \ell_1 < \dots < \ell_s < \ell_{s+1} = 1$. Let rewrite the vector of adaptable quantization levels as $\tilde{\ell} = [\tilde{\ell}_1, \dots, \tilde{\ell}_{2s+2}]^\top$ where $\tilde{\ell}_1 = -1 < \tilde{\ell}_2 < \dots < \tilde{\ell}_{2s+1} < \tilde{\ell}_{2s+2} = 1$. For $\theta \in [-1, 1]$, let $\tilde{\ell}(\theta)$ and $\rho(\theta)$ satisfy $\tilde{\ell}_{\tilde{\ell}(\theta)} \leq \theta \leq \tilde{\ell}_{\tilde{\ell}(\theta)+1}$ and $\theta = (1 - \rho(\theta))\tilde{\ell}_{\tilde{\ell}(\theta)} + \rho(\theta)\tilde{\ell}_{\tilde{\ell}(\theta)+1}$, respectively. Let $\mathbf{v} \in \mathbb{R}^d$ and $\theta_i = v_i / \|\mathbf{v}\|$ for $i = 1, \dots, d$.

Definition 1. The symmetrical quantization of a vector $\mathbf{v} \in \mathbb{R}^d$ is

$$Q_{\tilde{\ell}}(\mathbf{v}) \triangleq [q_{\tilde{\ell}}(v_1), \dots, q_{\tilde{\ell}}(v_d)]^\top, \quad (26)$$

where $q_{\tilde{\ell}}(v_i) = \|\mathbf{v}\| \cdot h(\theta_i)$ and $h(\theta_i)$'s are independent random variables such that $h(\theta_i) = \tilde{\ell}_{\tilde{\ell}(\theta_i)}$ with probability $1 - \rho(\theta_i)$ and $h(\theta_i) = \tilde{\ell}_{\tilde{\ell}(\theta_i)+1}$ otherwise where $\rho(\theta) = (\theta - \tilde{\ell}_{\tilde{\ell}(\theta)}) / (\tilde{\ell}_{\tilde{\ell}(\theta)+1} - \tilde{\ell}_{\tilde{\ell}(\theta)})$.

Let $r_i = |\theta_i|$ for $i = 1, \dots, d$. We have the following propositions.

Proposition 3. The variance of quantization with symmetric levels is given by

$$\mathbb{E}[\|Q_{\tilde{\ell}}(\mathbf{v}) - \mathbf{v}\|^2] = \|\mathbf{v}\|^2 \sum_{i=1}^d \sigma^2(r_i), \quad (27)$$

where

$$\sigma^2(r_i) = \sum_{r_i \in [0, \ell_1]} (\ell_1^2 - r_i^2) + \sum_{j=1}^s \sum_{r_i \in [\ell_j, \ell_{j+1}]} (r_i - \ell_j)(\ell_{j+1} - r_i). \quad (28)$$

Proof. Note that for symmetrical levels, we have

$$\begin{aligned} \sum_{|\theta_i| \in [\ell_j, \ell_{j+1}]} (|\theta_i| - \ell_j)(\ell_{j+1} - |\theta_i|) &= \sum_{\theta_i \in [-\ell_{j+1}, -\ell_j]} (\theta_i + \ell_{j+1})(-\ell_j - \theta_i) \\ &\quad + \sum_{\theta_i \in [\ell_j, \ell_{j+1}]} (\theta_i - \ell_j)(\ell_{j+1} - \theta_i). \end{aligned}$$

In addition, we have

$$\sum_{|\theta_i| \in [0, \ell_1]} (\ell_1^2 - |\theta_i|^2) = \sum_{\theta_i \in [-\ell_1, \ell_1]} (\ell_1^2 - \theta_i^2).$$

□

Proposition 4. If PDF of normalized gradients is an even function, i.e., $p(-\theta) = p(\theta)$ for $\theta \in [-1, 1]$, the expected normalized variance in (3) can be rewritten as

$$\Psi(\ell) = 2 \left(\int_0^{\ell_1} (\ell_1^2 - r^2) dF(r) + \sum_{j=1}^s \int_{\ell_j}^{\ell_{j+1}} (\ell_{j+1} - r)(r - \ell_j) dF(r) \right). \quad (29)$$

B.3.1 GD

For the case of symmetrical levels, the gradient of the expected normalized variance is given by

$$\begin{aligned}\frac{1}{2} \frac{\partial \Psi(\ell(t))}{\partial \ell_1} &= 2\ell_1(t)(F(\ell_1(t)) - F(0)) - \int_{\ell_1(t)}^{\ell_2(t)} (\ell_2(t) - r) dF(r), \\ \frac{1}{2} \frac{\partial \Psi(\ell(t))}{\partial \ell_j} &= \int_{\ell_{j-1}(t)}^{\ell_j(t)} (r - \ell_{j-1}(t)) dF(r) - \int_{\ell_j(t)}^{\ell_{j+1}(t)} (\ell_{j+1}(t) - r) dF(r)\end{aligned}\quad (30)$$

for $j = 2, \dots, s$.

For the special case of normal or truncated normal distribution, $\frac{1}{2} \frac{\partial \Psi(\ell(t))}{\partial \ell_j}$ is obtained by the R.H.S. of Eq. (25) for $j = 2, \dots, s$. In addition, we have:

$$\frac{1}{2} \frac{\partial \Psi(\ell(t))}{\partial \ell_1} = 2\ell_1(t)(F(\ell_1(t)) - F(0)) + (\mu - \ell_2(t))(F(\ell_2(t)) - F(\ell_1(t))) - \sigma^2 p(\ell_2(t)).$$

B.3.2 CD

In the following lemma, we solve the problem of optimizing a single level, *i.e.*, $\min_b \tilde{Q}(b)$ where

$$\tilde{Q}(b) = \int_0^b (b^2 - r^2) dF(r) + \int_b^c (c - r)(r - b) dF(r).$$

Proposition 5. *The optimal solution to minimize \tilde{Q} satisfies*

$$2b^*(F(b^*) - F(0)) = \int_{b^*}^c (c - r) dF(r), \quad (31)$$

where F is the CDF of the normalized coordinate.

Proof. Using Leibniz's rule [26], we have

$$\begin{aligned}\frac{d\tilde{Q}}{db} &= \int_0^b 2b dF(r) - \int_b^c (c - r) dF(r), \\ \frac{d^2\tilde{Q}}{db^2} &= bp(b) + cp(b) + 2(F(b) - F(0)) > 0.\end{aligned}$$

We note that \tilde{Q} is convex so we can find the unique optimal solution by satisfying the first order optimality condition. \square

We can solve Eq. (31) efficiently through a bisection search on $[0, \ell_2(t)]$. In particular, for the special case of normal and truncated normal, starting with a random $\ell(0)$, the update rule at iteration $t + 1$ is the same as Eq. (5) for $j = 2, \dots, s$. For $\ell_1(t + 1)$, we solve

$$\begin{aligned}(\ell_2(t) - \mu)(F(\ell_2(t)) - F(\ell_1(t + 1))) + 2\ell_1(t + 1)(F(0) - F(\ell_1(t + 1))) \\ + \sigma^2(p(\ell_2(t)) - p(\ell_1(t + 1))) = 0.\end{aligned}$$

B.3.3 Exponentially spaced levels

In this section, we focus on $\ell = [-1, -p, \dots, -p^s, p^s, \dots, p, 1]^\top$, *i.e.*, exponentially spaced levels with symmetry. Following Proposition 4, the expected normalized variance is given by

$$\Psi(p) = 2 \left(\int_0^{p^s} (p^{2s} - r^2) dF(r) + \sum_{j=0}^{s-1} \int_{p^{j+1}}^{p^j} (p^j - r)(r - p^{j+1}) dF(r) \right). \quad (32)$$

Using Leibniz's rule, we can compute the first order derivative:

$$\frac{1}{2} \frac{d\Psi(p)}{dp} = \int_0^{p^s} 2sp^{2s-1} dF(r) + \sum_{j=0}^{s-1} \int_{p^{j+1}}^{p^j} ((jp^{j-1} + (j+1)p^j)r - (2j+1)p^{2j}) dF(r).$$

In particular, in the special case of a normal or truncated normal distribution, we have

$$\begin{aligned} \frac{1}{2} \frac{d\Psi(p)}{dp} &= 2sp^{2s-1} (F(p^s) - F(0)) + \sigma^2 \sum_{j=0}^{s-1} (jp^{j-1} + (j+1)p^j) (p(p^{j+1}) - p(p^j)) \\ &\quad + \sum_{j=0}^{s-1} (\mu(jp^{j-1} + (j+1)p^j) - (2j+1)p^{2j}) (F(p^j) - F(p^{j+1})). \end{aligned}$$

We can update p efficiently by a gradient descent algorithm as we have a closed-form expression to find the gradient function.

C Expected variance minimization in Section 3.4

In the following, we provide the update rules and the analysis of computation complexity of ALQ, GD, and AMQ.

C.1 ALQ (CD update)

Starting with a random $\ell(0)$, for $t = 0, 1, \dots$ and $j = 1, \dots, s$, we solve

$$\bar{F}(\ell_j(t+1)) = \bar{F}(\ell_{j+1}(t)) - \int_{\ell_{j-1}(t)}^{\ell_{j+1}(t)} \frac{r - \ell_{j-1}(t)}{\ell_{j+1}(t) - \ell_{j-1}(t)} d\bar{F}(r) \quad (33)$$

by a bisection search on $[\ell_{j-1}(t), \ell_{j+1}(t)]$.

In the special case of (truncated) normal distribution, to obtain $\ell_j(t+1)$ for $j = 1, \dots, s$, we solve

$$\begin{aligned} \sum_{n=1}^N \gamma_n ((\ell_{j-1}(t) - \mu_n) (F_n(\ell_{j+1}(t)) - F_n(\ell_{j-1}(t))) + \sigma_n^2 (p_n(\ell_{j+1}(t)) - p_n(\ell_{j-1}(t)))) \\ + (\ell_{j+1}(t) - \ell_{j-1}(t)) (\bar{F}(\ell_{j+1}(t)) - \bar{F}(\ell_j(t+1))) = 0 \end{aligned} \quad (34)$$

by a bisection search on $[\ell_{j-1}(t), \ell_{j+1}(t)]$.

In the special case of symmetrical levels and (truncated) normal distribution, the update rule is the same as Eq. (34) for $j = 2, \dots, s$. For $\ell_1(t+1)$, we solve

$$\begin{aligned} \sum_{n=1}^N \gamma_n ((\ell_2(t) - \mu_n) (F_n(\ell_2(t)) - F_n(\ell_1(t+1))) + \sigma_n^2 (p_n(\ell_2(t)) - p_n(\ell_1(t+1)))) \\ + 2\ell_1(t+1) (\bar{F}(0) - \bar{F}(\ell_1(t+1))) = 0 \end{aligned} \quad (35)$$

by a bisection search on $[0, \ell_2(t)]$.

C.2 GD update

The GD algorithm to minimize Eq. (10) is based on the following update rule by starting from a random $\ell(0)$:

$$\begin{aligned} \ell_j(t+1) &= \ell_j(t) - \text{sign}(\hat{g}(t, j)) \min\{\eta(t)|\hat{g}(t, j)|, \delta_j(t)/2\} \\ \hat{g}(t, j) &= \int_{\ell_{j-1}(t)}^{\ell_j(t)} (r - \ell_{j-1}(t)) d\bar{F}(r) - \int_{\ell_j(t)}^{\ell_{j+1}(t)} (\ell_{j+1}(t) - r) d\bar{F}(r) \end{aligned} \quad (36)$$

for $t = 0, 1, \dots$ and $j = 1, \dots, s$.

In the special case of (truncated) normal distribution, we have

$$\begin{aligned} \hat{g}(t, j) = & \sum_{n=1}^N \gamma_n \left((\mu_n - \ell_{j-1}(t)) (F_n(\ell_j(t)) - F_n(\ell_{j-1}(t))) \right. \\ & \left. + (\mu_n - \ell_{j+1}(t)) (F_n(\ell_{j+1}(t)) - F_n(\ell_j(t))) + \sigma_n^2 (p_n(\ell_{j-1}(t)) - p_n(\ell_{j+1}(t))) \right). \end{aligned} \quad (37)$$

C.3 AMQ (GD update with exponentially spaced levels)

In this section, we focus on $\ell = [-1, -p, \dots, -p^s, p^s, \dots, p, 1]^\top$, *i.e.*, exponentially spaced levels with symmetry. Following Proposition 4, the expected variance of quantization is given by

$$\tilde{V}(p) = 2 \left(\int_0^{p^s} (p^{2s} - r^2) d\bar{F}(r) + \sum_{j=0}^{s-1} \int_{p^{j+1}}^{p^j} (p^j - r)(r - p^{j+1}) d\bar{F}(r) \right). \quad (38)$$

In the special case of (truncated) normal distribution, we have

$$\begin{aligned} \frac{d\tilde{V}(p)}{dp} = & +2sp^{2s-1} (\bar{F}(p^s) - \bar{F}(0)) \\ & + \sum_{n=1}^N \gamma_n \left(\sum_{j=0}^{s-1} (\mu_n(jp^{j-1} + (j+1)p^j) - (2j+1)p^{2j}) (F_n(p^j) - F_n(p^{j+1})) \right. \\ & \left. + \sigma_n^2 \sum_{j=0}^{s-1} (jp^{j-1} + (j+1)p^j) (p_n(p^{j+1}) - p_n(p^j)) \right). \end{aligned}$$

We can update p efficiently by a gradient descent algorithm as we have a closed-form expression to find the gradient function.

C.4 Computational complexity and scalability

The number of iterations for ALQ method to converge is in the order of $O(s \log(1/\epsilon))$ where ϵ is the suboptimality gap of bisection search. The number of iterations for AMQ method to achieve a local minimum with gap ϵ is $O(1/\epsilon)$. The total number of gradient computations for GD method to achieve a local minimum with gap ϵ is $O(s/\epsilon)$. Note that processors can run our methods in parallel. The time complexity of these methods is independent of the number of samples, the number of processors, and the number of parameters. The extra computational overhead is negligible compared to costs of computation of stochastic gradients and communication. Furthermore, we do not need to optimize levels at each iteration. Our experimental results suggest that it is sufficient to optimize levels at the `lr_scheduler` iterations.

D Encoding

A quantized gradient $Q_\ell(\mathbf{v})$ can be uniquely determined by a tuple $(\|\mathbf{v}\|, \mathbf{s}, \mathbf{h})$ where $\|\mathbf{v}\|$ is the Euclidean norm of the gradient, $\mathbf{s} \triangleq [\text{sign}(v_1), \dots, \text{sign}(v_d)]^\top$ is the vector of signs of the coordinates v_i 's, and $\mathbf{h} \triangleq [h(r_1), \dots, h(r_d)]^\top$ are the discrete values of the normalized coordinates after quantization.

We can describe the ENCODE function (for Algorithm 1) in terms of the tuple $(\|\mathbf{v}\|, \mathbf{s}, \mathbf{h})$ and encoding/decoding scheme $\Gamma : \{\ell_0, \ell_1, \dots, \ell_{s+1}\} \rightarrow \{0, 1\}^*$ and $\Gamma^{-1} : \{0, 1\}^* \rightarrow \{\ell_0, \ell_1, \dots, \ell_{s+1}\}$. Any lossless prefix code can be used for encoding/decoding. In particular, we consider Huffman coding due to its efficient encoding/decoding and its optimality in terms of achieving the minimum expected code length among methods encoding symbols separately [32].

The encoding, $\text{ENCODE}(\mathbf{v})$, of a stochastic gradient is as follows: We first encode the norm $\|\mathbf{v}\|$ using b bits where, in practice, we use standard 32-bit floating point encoding. We then proceed in rounds, $t = 0, 1, \dots, d$. Noting that we do not need to encode the sign bit for zero entries of \mathbf{h} , on

round t , if $h(r_t) = 0$, we transmit $\Gamma(0)$. If $h(r_t) \neq 0$, we transmit $\Gamma(h_{r_t})$, transmit one bit encoding the sign(v_t), and proceed to the next entry of \mathbf{h} .

The DECODE function (for Algorithm 1) simply reads b bits to reconstruct $\|\mathbf{v}\|$. Using Γ^{-1} , it decodes the index of the first coordinate, depending on whether the decoded entry is zero or nonzero, it may read one bit indicating the sign, and then proceeds to decode the next symbol. The process proceeds in rounds, mimicking the encoding process, finishing when all coordinates have been decoded. Note that we can improve coding efficiency by encoding blocks of symbols at the cost of increasing encoding/decoding complexity. In this paper, we focus on a simple lossless prefix coding scheme that encodes symbols separately.

In order to implement an efficient lossless prefix code, we need to know the probabilities associated with our symbols to be coded, *i.e.*, $\{\ell_0, \ell_1, \dots, \ell_{s+1}\}$. Fortunately, we can compute those probabilities using the marginal PDF of normalized coordinates and quantization levels as shown in this proposition:

Proposition 6. *The probability of occurrence of ℓ_j (weight of symbol ℓ_j) is given by*

$$\Pr(\ell_j) = \int_{\ell_{j-1}}^{\ell_j} \frac{r - \ell_{j-1}}{\ell_j - \ell_{j-1}} dF(r) + \int_{\ell_j}^{\ell_{j+1}} \frac{\ell_{j+1} - r}{\ell_{j+1} - \ell_j} dF(r)$$

for $j = 1, \dots, s$ where F is the marginal CDF of normalized coordinates. In addition, we have

$$\Pr(\ell_0 = 0) = \int_0^{\ell_1} \frac{1 - r}{\ell_1} dF(r) \text{ and } \Pr(\ell_{s+1} = 1) = \int_{\ell_s}^1 \frac{r - \ell_s}{1 - \ell_s} dF(r).$$

In the special case of truncated normal distribution, we have the symbol probabilities in closed-form:

Corollary 3. *Suppose normalized coordinates have truncated normal distribution with PDF $p_{\mathcal{T}}$ and CDF $F_{\mathcal{T}}$ defined in Appendix A.2. The probability of occurrence of ℓ_j (weight of symbol ℓ_j) is given by*

$$\begin{aligned} \Pr(\ell_j) = & \frac{\sigma^2(p_{\mathcal{T}}(\ell_{j-1}) - p_{\mathcal{T}}(\ell_j)) + (\mu - \ell_{j-1})(F_{\mathcal{T}}(\ell_j) - F_{\mathcal{T}}(\ell_{j-1}))}{\ell_j - \ell_{j-1}} \\ & + \frac{\sigma^2(p_{\mathcal{T}}(\ell_{j+1}) - p_{\mathcal{T}}(\ell_j)) + (\ell_{j+1} - \mu)(F_{\mathcal{T}}(\ell_{j+1}) - F_{\mathcal{T}}(\ell_j))}{\ell_{j+1} - \ell_j} \end{aligned}$$

for $j = 1, \dots, s$. In addition, we have

$$\begin{aligned} \Pr(\ell_0 = 0) &= \frac{\sigma^2(p_{\mathcal{T}}(\ell_1) - p_{\mathcal{T}}(0)) + (\ell_1 - \mu)(F_{\mathcal{T}}(\ell_1) - F_{\mathcal{T}}(0))}{\ell_1}, \\ \Pr(\ell_{s+1} = 1) &= \frac{\sigma^2(p_{\mathcal{T}}(\ell_s) - p_{\mathcal{T}}(1)) + (\mu - \ell_s)(F_{\mathcal{T}}(1) - F_{\mathcal{T}}(\ell_s))}{1 - \ell_s}. \end{aligned}$$

Note that each processor can construct the Huffman tree by knowing ℓ and estimating μ and σ^2 . A Huffman tree of a source with n symbols can be constructed in time $O(n)$ if the symbols are sorted by probability. Huffman codes are optimal in terms of expected code-length:

Theorem 5 (Cover and Thomas 32, Theorems 5.4.1 and 5.8.1). *Let X denote a random source. The expected code-length of an optimal prefix code, *e.g.*, Huffman code to compress X is bounded by $H(X) \leq \mathbb{E}[L] \leq H(X) + 1$ where $H(X)$ is the entropy of X in bits.*

E Variance gap

Proposition 7 (Variance gap). *For any distribution where the gap between the expected variance of a normalized coordinate under an optimal quantization to minimize (3) and the worst-case one is lower bounded by some constant, the total gap is lower bounded by $\Omega(d)$. We quantify this gap for the special case of one level with truncated normal density.*

Proof. Suppose we want to design a single level $b \in (0, 1)$. As shown in Corollary 2, the optimal b to minimize $Q(b)$ is given by $b^* = F^{-1}(1 - \mathbb{E}[R])$. Let R have truncated normal density with

parameters μ, σ^2 in the unit interval. Plugging PDF and CDF of R , the optimal level to minimize (3) is given by

$$b^* = \sigma \Phi^{-1} \left(\Delta(1 - \mu) + \sigma \delta + \Phi \left(-\frac{\mu}{\sigma} \right) \right) + \mu,$$

where $\Delta = \Phi((1 - \mu)/\sigma) - \Phi(-\mu/\sigma)$, $\delta = \phi((1 - \mu)/\sigma) - \phi(-\mu/\sigma)$, $\Phi(x) = \int_{-\infty}^x \exp(-u^2/2) du / \sqrt{2\pi}$, and $\phi(x) = \exp(-x^2/2) / \sqrt{2\pi}$.

Note that $\hat{b} = 1/2$ minimizes the worst-case variance upper bound in Eq. (1) [27]. In general, $b^* \neq \hat{b}$ depending on μ and σ^2 . Without loss of generality, assume $b^* > \hat{b}$.

As show in Proposition 2, $Q(b) = \int_0^b (b - r)(r - a) dF(r) + \int_b^1 (1 - r)(r - b) dF(r)$ is convex and

$$\frac{d^2 Q}{db^2} = \frac{\phi((b - \mu)/\sigma)}{\sigma \Delta}.$$

In the interval $[\hat{b}, b^*]$, we have

$$\frac{d^2 Q}{db^2} \geq \gamma \triangleq \min \left\{ \frac{\phi((b^* - \mu)/\sigma)}{\sigma \Delta}, \frac{\phi((\hat{b} - \mu)/\sigma)}{\sigma \Delta} \right\}.$$

In this interval, Q is γ -strongly convex, i.e.,

$$Q(\hat{b}) \geq Q(b^*) + \frac{\gamma}{2}(b^* - \hat{b})^2.$$

Hence, the gap in the expected normalized variance under b^* and \hat{b} is lower bounded by:

$$\frac{1}{2} d\gamma(b^* - \hat{b})^2.$$

□

F Proof of Theorem 2 (variance bound)

Let $\ell = [\ell_0, \ell_1, \dots, \ell_s, \ell_{s+1}]^\top$ denote arbitrary quantization levels where $\ell_0 = 0 < \ell_1 < \dots < \ell_{s+1} = 1$. The variance of $Q_\ell(\mathbf{v})$, i.e., $V_\ell(\mathbf{v}) = \mathbb{E}[\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2]$, can be expressed as

$$V_\ell(\mathbf{v}) = \|\mathbf{v}\|_q^2 \left(\sum_{r_i \in \mathcal{I}_0} (\ell_1 - r_i) r_i + \sum_{j=1}^s \sum_{r_i \in \mathcal{I}_j} (\ell_{j+1} - r_i)(r_i - \ell_j) \right), \quad (39)$$

where $r_i = |v_i| / \|\mathbf{v}\|_q$, $\mathcal{I}_0 \triangleq [0, \ell_1]$, and $\mathcal{I}_j \triangleq [\ell_j, \ell_{j+1}]$ for $j = 1, \dots, s$.

We can find the minimum k_j that satisfies $(\ell_{j+1} - r)(r - \ell_j) \leq k_j r^2$ for $r \in \mathcal{I}_j$ and $j = 1, \dots, s$. Expressing $r = \ell_j \theta$, we can find k through solving

$$\begin{aligned} k_j &= \max_{1 \leq \theta \leq \ell_{j+1}/\ell_j} \frac{(\ell_{j+1}/\ell_j - \theta)(\theta - 1)}{\theta^2} \\ &= \frac{(\ell_{j+1}/\ell_j - 1)^2}{4(\ell_{j+1}/\ell_j)}. \end{aligned} \quad (40)$$

We note that $\ell_{j+1}/\ell_j > 1$ and $(x - 1)^2/(4x)$ is monotonically increasing function of x for $x > 1$.

Furthermore, note that

$$\sum_{r_i \notin \mathcal{I}_0} r_i^2 \leq \frac{\|\mathbf{v}\|_2^2}{\|\mathbf{v}\|_q^2}.$$

Substituting Eq. (40) into Eq. (39), an upper bound on $V_\ell(\mathbf{v})$ is given by

$$V_\ell(\mathbf{v}) \leq \|\mathbf{v}\|_q^2 \left(\frac{(\ell_{j^*+1}/\ell_{j^*} - 1)^2}{4(\ell_{j^*+1}/\ell_{j^*})} \frac{\|\mathbf{v}\|_2^2}{\|\mathbf{v}\|_q^2} + \sum_{r_i \in \mathcal{I}_0} (\ell_1 - r_i) r_i \right),$$

where $j^* = \arg \max_{1 \leq j \leq s} \ell_{j+1}/\ell_j$.

In our proofs, we use the following known lemma.

Lemma 1. Let $\mathbf{v} \in \mathbb{R}^d$. Then, for all $0 < p < q$, we have $\|\mathbf{v}\|_q \leq \|\mathbf{v}\|_p \leq d^{1/p-1/q} \|\mathbf{v}\|_q$.

Note that Lemma 1 holds even when $q < 1$ and $\|\cdot\|_q$ is merely a seminorm.

In the following, we derive a bound on $\sum_{r_i \in \mathcal{I}_0} (2^{-s} - r_i) r_i$, which completes the proof.

Lemma 2. Let $p \in (0, 1)$ and $r \in \mathcal{I}_0$. Then we have $r(\ell_1 - r) \leq K_p \ell_1^{(2-p)} r^p$ where

$$K_p = \left(\frac{1/p}{2/p-1} \right) \left(\frac{1/p-1}{2/p-1} \right)^{(1-p)}. \quad (41)$$

Proof. We can find the minimum K_p through solving $K_p = \ell_1^{(-2+p)} \max_{r \in \mathcal{I}_0} r(\ell_1 - r)/r^p$. Expressing the optimization variable as $r = \ell_1 \theta^{1/p}$, K_p can be obtained by solving this problem:

$$K_p = \max_{0 < \theta < 1} \theta^{1/p-1} - \theta^{2/p-1}. \quad (42)$$

We can solve (42) and obtain the optimal solution $\theta^* = \left(\frac{1/p-1}{2/p-1} \right)^p$. Substituting θ^* into (42), we obtain Eq. (41). \square

Let \mathcal{S}_j denote the coordinates of vector \mathbf{v} whose elements fall into the $(j+1)$ -th bin, i.e., $\mathcal{S}_j \triangleq \{i : r_i \in [l_j, l_{j+1}]\}$ for $j = 0, \dots, s$.

Then, for any $0 < p < 1$ and $q \geq 2$, we have

$$\begin{aligned} \|\mathbf{v}\|_q^2 \sum_{r_i \in \mathcal{I}_0} r_i^p &= \|\mathbf{v}\|_q^{2-p} \sum_{i \in \mathcal{S}_0} |v_i|^p \\ &\leq \|\mathbf{v}\|_q^{2-p} \|\mathbf{v}\|_p^p \\ &\leq \|\mathbf{v}\|_q^{2-p} \|\mathbf{v}\|_2^p d^{1-p/2} \\ &\leq \|\mathbf{v}\|_2^2 d^{1-p/2}, \end{aligned}$$

where the third inequality holds as $\|\mathbf{v}\|_p \leq \|\mathbf{v}\|_2 d^{1/p-1/2}$ using Lemma 1 and the last inequality holds as $\|\mathbf{v}\|_q \leq \|\mathbf{v}\|_2$ for $q \geq 2$.

This gives us an upper bound on $V_\ell(\mathbf{v})$:

$$V_\ell(\mathbf{v}) \leq \|\mathbf{v}\|_2^2 \left(\frac{(\ell_{j^*+1}/\ell_{j^*} - 1)^2}{4(\ell_{j^*+1}/\ell_{j^*})} + K_p \ell_1^{(2-p)} d^{1-p/2} \right).$$

For $q \geq 1$, we have $\|\mathbf{v}\|_q^{2-p} \leq \|\mathbf{v}\|_2^{2-p} d^{\frac{2-p}{\min\{q, 2\}} - \frac{2-p}{2}}$, which gives

$$V_\ell(\mathbf{v}) \leq \|\mathbf{v}\|_2^2 \left(\frac{(\ell_{j^*+1}/\ell_{j^*} - 1)^2}{4(\ell_{j^*+1}/\ell_{j^*})} + K_p \ell_1^{(2-p)} d^{\frac{2-p}{\min\{q, 2\}}} \right).$$

G Proof of Theorem 3 (code-length bound)

Let $|\cdot|$ denote the length of a binary string. In this section, we obtain an upper bound on $\mathbb{E}[|\text{ENCODE}(\mathbf{v})|]$, i.e., the expected number of communication bits per iteration. Recall from Appendix D that the quantized vectors $Q_\ell(\mathbf{v})$ is uniquely determined by the tuple $(\|\mathbf{v}\|_q, \mathbf{s}, \mathbf{h})$.

We first encode the norm $\|\mathbf{v}\|_q$ using b bits where, in practice, we use standard 32-bit floating point encoding.

We send one bit for each nonzero entry of \mathbf{h} . Let $\mathcal{S}_j \triangleq \{i : r_i \in [l_j, l_{j+1}]\}$ and $d_j \triangleq |\mathcal{S}_j|$ for $j = 0, \dots, s$. We have an upper bound on the expected number of nonzero entries as follows:

Lemma 3. Let $\mathbf{v} \in \mathbb{R}^d$. The expected number of nonzeros in $Q_\ell(\mathbf{v})$ is bounded above by

$$\mathbb{E}[|Q_\ell(\mathbf{v})|_0] \leq \ell_1^{-q} + \frac{d^{1-1/q}}{\ell_1}.$$

Proof. Note that $d - d_0 \leq \ell_1^{-q}$ since

$$(d - d_0)\ell_1^q \leq \sum_{i \notin \mathcal{S}_0} r_i^q \leq 1. \quad (43)$$

For each $i \in \mathcal{S}_0$, $Q_\ell(v_i)$ becomes zero with probability $1 - r_i/\ell_1$, which results in

$$\begin{aligned} \mathbb{E}[\|Q_\ell(\mathbf{v})\|_0] &\leq d - d_0 + \sum_{i \in \mathcal{S}_0} r_i/\ell_1 \\ &\leq \ell_1^{-q} + \frac{d^{1-1/q}}{\ell_1}, \end{aligned} \quad (44)$$

where the last inequality holds as $\|\mathbf{v}\|_1 \leq \|\mathbf{v}\|_q d^{1-1/q}$ using Lemma 1. \square

For each entry of \mathbf{h} , we send the associated codeword. The optimal expected code-length for transmitting one random symbol is within one bit of the entropy of the source. Hence, we need to transmit upto $d(H(L) + 1)$ to transmit entries of \mathbf{h} [32]. Putting everything together, we have

$$\mathbb{E}[\|\text{ENCODE}(\mathbf{v})\|] \leq b + n_{\ell_1, d} + d(H(L) + 1).$$

Finally, note that the entropy of a source with n outcomes is bounded above by $\log_2(n)$.

H AQSGD for smooth nonconvex optimization

On nonconvex problems, we can establish convergence guarantees in terms of convergence to a local minima for a smooth loss function along the lines of, e.g., [33, Theorem 2.1].

Theorem 6 (AQSGD for smooth nonconvex optimization). *Let $f : \Omega \rightarrow \mathbb{R}$ denote a possibly nonconvex and β -smooth function. Let $\mathbf{w}_0 \in \Omega$ denote an initial point, $\overline{\epsilon_Q}$ and $\overline{N_Q}$ be defined as in Theorem 4, $T \in \mathbb{Z}^{>0}$, and $f^* = \inf_{\mathbf{w} \in \Omega} f(\mathbf{w})$. Suppose that Algorithm 1 is executed for T iterations with a learning rate $\alpha < 2/\beta$ on M processors, each with access to independent stochastic gradients of f with a second-moment bound B , such that levels are updated K times where ℓ_k with variance bound $\epsilon_{Q,k}$ and code-length bound $N_{Q,k}$ is used for T_k iterations. Then there exists a random stopping time $R \in \{0, \dots, T\}$ such that AQSGD guarantees*

$$\mathbb{E}[\|\nabla f(\mathbf{w}_R)\|^2] \leq \frac{\beta(f(\mathbf{w}_0) - f^*)}{T} + \frac{2(1 + \overline{\epsilon_Q})B}{M}.$$

In addition, AQSGD requires at most $\overline{N_Q}$ communication bits per iteration in expectation.

I AQSGD with momentum

The update rule for full-precision unified momentum SGD (UMSGD) is given by [34]

$$\begin{aligned} \mathbf{y}_{t+1} &= \mathbf{w}_t - \alpha g(\mathbf{w}_t) \\ \mathbf{y}_{t+1}^\ell &= \mathbf{w}_t - l\alpha g(\mathbf{w}_t) \\ \mathbf{w}_{t+1} &= \mathbf{y}_{t+1} + \mu(\mathbf{y}_{t+1}^\ell - \mathbf{y}_t^\ell), \end{aligned} \quad (45)$$

where \mathbf{w}_t is the current parameter input and $\mu \in [0, 1)$ is the momentum parameter. Note that the heavy-ball method [35] and Nesterov's accelerated gradient method [36] are the special cases of UMSGD obtained by substituting $l = 0$ and $l = 1$ into Eq. (45), respectively.

The steps for data-parallel version of UMSGD are those in Algorithm 1 by replacing Line 9 with an UMSGD update. We have convergence guarantees for adaptively quantized SGD with momentum (AQSGDM) along the lines of, e.g., [34, Theorem 1]. We first establish the convergence guarantees for convex optimization in the following theorem.

Theorem 7 (AQSGDM for convex optimization). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ denote a convex function with $\|\nabla f(\mathbf{w})\| \leq V$ for all \mathbf{w} . Let \mathbf{w}_0 denote an initial point, $\mathbf{w}^* = \arg \min f(\mathbf{w})$, $\hat{\mathbf{w}}_T = 1/T \sum_{t=0}^T \mathbf{w}_t$, and $\overline{\epsilon_Q}$ and $\overline{N_Q}$ be defined as in Theorem 4.*

Suppose that AQSGDM is executed for T iterations with a learning rate $\alpha > 0$ on M processors, each with access to independent stochastic gradients of f with a second-moment bound B , such that levels are updated K times where ℓ_k with variance bound $\epsilon_{Q,k}$ and code-length bound $N_{Q,k}$ is used for T_k iterations. Then AQSGDM satisfies

$$\mathbb{E}[f(\hat{\mathbf{w}}_T)] - \min_{\mathbf{w} \in \Omega} f(\mathbf{w}) \leq \epsilon_\mu^\ell, \quad (46)$$

where $\epsilon_\mu^\ell = \mu(f(\mathbf{w}_0) - f(\mathbf{w}^*)) / ((1 - \mu)(T + 1)) + (1 - \mu)\|\mathbf{w}_0 - \mathbf{w}^*\|^2 / (2\alpha(T + 1)) + \alpha(1 + 2l\mu)(V^2 + (1 + \bar{\epsilon}_Q)B/M) / (2(1 - \mu))$.

In addition, AQSGD requires at most \bar{N}_Q communication bits per iteration in expectation.

On nonconvex problems, (weaker) convergence guarantees can be established for AQSGDM. In particular, AQSGDM is guaranteed to converge to a local minima for smooth general loss functions.

Theorem 8 (AQSGDM for smooth nonconvex optimization). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ denote a possibly nonconvex and β -smooth function with $\|\nabla f(\mathbf{w})\| \leq V$ for all \mathbf{w} . Let \mathbf{w}_0 denote an initial point, $\mathbf{w}^* = \arg \min f(\mathbf{w})$, and $\bar{\epsilon}_Q$ and \bar{N}_Q be defined as in Theorem 4.*

Suppose that AQSGDM is executed for T iterations with $\alpha = \min\{(1 - \mu)/(2\beta), C/\sqrt{T + 1}\}$ for some $C > 0$ on M processors, each with access to independent stochastic gradients of f with a second-moment bound B , such that levels are updated K times where ℓ_k with variance bound $\epsilon_{Q,k}$ and code-length bound $N_{Q,k}$ is used for T_k iterations. Then AQSGDM satisfies

$$\min_{t=0,\dots,T} \mathbb{E}[\|\nabla f(\mathbf{w}_t)\|^2] \leq \frac{2(f(\mathbf{w}_0) - f(\mathbf{w}^*))(1 - \mu)}{\alpha(T + 1)} + \frac{C\tilde{V}}{(1 - \mu)^3\sqrt{T + 1}},$$

where

$$\tilde{V} = \beta(\mu^2((1 - \mu)l - 1)^2 + (1 - \mu)^2)(V^2 + (1 + \bar{\epsilon}_Q)B/M).$$

In addition, AQSGD requires at most \bar{N}_Q communication bits per iteration in expectation.

J Theoretical guarantees for levels with symmetry

We first obtain variance upper bound for the symmetrical $\ell = [-\ell_{s+1}, \dots, -\ell_1, \ell_1, \dots, \ell_{s+1}]^\top$.

Theorem 9 (Variance bound). *Let $\mathbf{v} \in \mathbb{R}^d$ and $q \geq 1$. The quantization of \mathbf{v} under L^q normalization satisfies $\mathbb{E}[Q_\ell(\mathbf{v})] = \mathbf{v}$. Furthermore, we have*

$$\mathbb{E}[\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2] \leq \epsilon_Q \|\mathbf{v}\|_2^2, \quad (47)$$

where $\epsilon_Q = \ell_1^2 d^{\frac{2}{\min\{q, 2\}}} + \frac{(\ell_{j^*+1}/\ell_{j^*}-1)^2}{4(\ell_{j^*+1}/\ell_{j^*})}$ where $j^* = \arg \max_{1 \leq j \leq s} \ell_{j+1}/\ell_j$.

Proof. Following Proposition 3, the variance is given by

$$\mathbb{E}[\|Q_\ell(\mathbf{v}) - \mathbf{v}\|_2^2] = \|\mathbf{v}\|_q^2 \left(\sum_{r_i \in [0, \ell_1]} (\ell_1^2 - r_i^2) + \sum_{j=1}^s \sum_{r_i \in [\ell_j, \ell_{j+1}]} (r_i - \ell_j)(\ell_{j+1} - r_i) \right).$$

Note that $\ell_1^2 - r^2 \leq \ell_1^2$ for $r \in [0, \ell_1]$. The rest of the proof follows the proof of Theorem 2. \square

In order to implement an efficient lossless prefix code, we need to know the probabilities associated with our symbols to be coded, i.e., $\{-\ell_{s+1}, \dots, -\ell_1, \ell_1, \dots, \ell_{s+1}\}$. We can obtain those probabilities using the marginal PDF of normalized coordinates:

Proposition 8. *Suppose $p(-\theta) = p(\theta)$. The probability of occurrence of ℓ_j (weight of symbol ℓ_j) $\Pr(\ell_j)$ is equal to $\Pr(-\ell_j)$, given by*

$$\Pr(\ell_j) = \int_{\ell_{j-1}}^{\ell_j} \frac{\theta - \ell_{j-1}}{\ell_j - \ell_{j-1}} dF(\theta) + \int_{\ell_j}^{\ell_{j+1}} \frac{\ell_{j+1} - \theta}{\ell_{j+1} - \ell_j} dF(\theta)$$

for $j = 2, \dots, s$. In addition, we have

$$\begin{aligned}\Pr(\ell_1) &= \Pr(-\ell_1) = \int_{-\ell_1}^{\ell_1} \frac{\theta + \ell_1}{2\ell_1} dF(\theta) + \int_{\ell_1}^{\ell_2} \frac{\ell_2 - \theta}{\ell_2 - \ell_1} dF(\theta), \\ \Pr(\ell_{s+1} = 1) &= \Pr(-\ell_{s+1}) = \int_{\ell_s}^1 \frac{\theta - \ell_s}{1 - \ell_s} dF(\theta).\end{aligned}$$

In the special case of truncated normal distribution, we have the symbol probabilities in closed-form:

Corollary 4. Suppose normalized coordinates have truncated normal distribution with PDF $p_{\mathcal{T}}$ and CDF $F_{\mathcal{T}}$ defined in Appendix A.2. The probability of occurrence of ℓ_j (weight of symbol ℓ_j) is given by

$$\begin{aligned}\Pr(\ell_j) &= \Pr(-\ell_j) = \frac{\sigma^2(p_{\mathcal{T}}(\ell_{j-1}) - p_{\mathcal{T}}(\ell_j)) + (\mu - \ell_{j-1})(F_{\mathcal{T}}(\ell_j) - F_{\mathcal{T}}(\ell_{j-1}))}{\ell_j - \ell_{j-1}} \\ &\quad + \frac{\sigma^2(p_{\mathcal{T}}(\ell_{j+1}) - p_{\mathcal{T}}(\ell_j)) + (\ell_{j+1} - \mu)(F_{\mathcal{T}}(\ell_{j+1}) - F_{\mathcal{T}}(\ell_j))}{\ell_{j+1} - \ell_j}\end{aligned}$$

for $j = 2, \dots, s$. In addition, we have

$$\begin{aligned}\Pr(\ell_1) &= \Pr(-\ell_1) = \frac{\sigma^2(p_{\mathcal{T}}(\ell_2) - p_{\mathcal{T}}(\ell_1)) + (\ell_2 - \mu)(F_{\mathcal{T}}(\ell_2) - F_{\mathcal{T}}(\ell_1))}{\ell_2 - \ell_1} \\ &\quad + \frac{(\mu + \ell_1)(F_{\mathcal{T}}(\ell_1) - F_{\mathcal{T}}(-\ell_1))}{2\ell_1}, \\ \Pr(\ell_{s+1} = 1) &= \Pr(-\ell_{s+1}) = \frac{\sigma^2(p_{\mathcal{T}}(\ell_s) - p_{\mathcal{T}}(1)) + (\mu - \ell_s)(F_{\mathcal{T}}(1) - F_{\mathcal{T}}(\ell_s))}{1 - \ell_s}.\end{aligned}$$

Finally, we have the following bound on the expected number of communication bits per iteration for quantizing with symmetrical levels.

Theorem 10 (Code-length bound). Let $\mathbf{v} \in \mathbb{R}^d$ and $q \geq 1$. The expectation $\mathbb{E}[|\text{ENCODE}(\mathbf{v})|]$ of the number of communication bits needed to transmit $Q_{\ell}(\mathbf{v})$ under L^q normalization is bounded by

$$\mathbb{E}[|\text{ENCODE}(\mathbf{v})|] \leq b + d(H(L) + 1) \leq b + d(\log_2(2s + 2) + 1), \quad (48)$$

where b is a constant and L is a random variable with the probability mass function given by Proposition 8.

K Experimental details and additional experiments

In this section, we provide additional experiments for the methods evaluated in Section 5. In addition to baselines discussed in Section 5, we present results for ALQ-N, which minimizes the expected normalized variance using coordinate descent in (3), ALQ with norm adjustments in Section 3.4, ALQ adapted using gradient descent in Section 3.2 (ALQG, ALQG-N), AMQ-N, and AMQ. This section includes full ImageNet runs. Figs. 3a, 3b, 5a and 5b have also been extended to include all variations of the proposed algorithms and baselines.

An implementation challenge is that the value of the statistics, especially the variance, can become very small. This makes PDF and CDF calculations challenging. The challenge is that the value of PDF is very close to zero when it is far from the mean but not exactly zero. In order to overcome this challenge, we use histograms to model the distribution of gradients as a weighted sum of truncated normals. Another problem is the large number of statistics that are calculated. As presented in Section 3, we sample a number of gradients and then normalize the gradients. Then we split the gradients into buckets and calculate average, variance, and norm of each of the buckets. The number of means, variances, and norms can become very large with large networks and small bucket sizes. To reduce computational complexity of the algorithm, we sample uniformly from these values. This number of samples is equal to 20 for small networks such as ResNet-8 and networks trained on CIFAR-10; however, in experiments on ImageNet, we used 350 samples to achieve the desired accuracy.

Table 3: Training Hyper-parameters for CIFAR-10 and ImageNet

Hyperparameter	ResNet-32 on CIFAR-10	ResNet-110 on CIFAR-10	ImageNet
Learning Rate	0.1	0.1	0.1
LR Decay Schedule	At 45K & 60K	At 45K & 60K	At 300K & 450K
Batch Size	128	64	64
Momentum	0.9	0.9	0.9
Total Iterations	80K	80K	600K
Weight Decay	10^{-4}	10^{-4}	10^{-4}
Optimizer	SGD	SGD	SGD

Table 4: Validation Accuracy on Full ImageNet Run

Quantization Method	ResNet-18 on ImageNet
Bucket Size	8192
SGD	$64.67\% \pm 0.13$
SuperSGD	$69.85\% \pm 0.05$
NUQSGD [21, 22]	$35.43\% \pm 0.28$
QSGDinf [20]	$67.48\% \pm 0.08$
TRN [15]	$63.97\% \pm 0.11$
ALQ	$68.65\% \pm 0.10$
ALQ-N	$68.50\% \pm 0.10$
AMQ	$67.76\% \pm 0.09$
AMQ-N	$67.96\% \pm 0.10$

One other understudied detail in quantizing is how bucketing is performed. In [20, 21], gradient coordinates in each bucket do not exceed the layer size. It means that the gradient coordinates in a bucket do not contain gradient coordinates from the next layer even if the bucket size is not fully utilized. This leads to creation of under-sized buckets that can be problematic for quantization performance. Different tricks are employed to fix this problem. These tricks include transmitting biases or under-sized buckets in full-precision (not that typically biases are main sources of under-sized buckets). In our implementation, we normalize the buckets network-wise and do not consider the layer size as the bucket size boundary. We only transmit the last bucket in full precision if it is smaller than the specified bucket size.

Update Schedule. In the ImageNet and CIFAR-10 runs, adaptive level updates are scheduled at 100 and 2000 iterations only once and every 10K iterations. The reason for this schedule is changes in the gradient statistics over the course of training. As shown in Fig. 1, the average variance changes rapidly during the first iterations and then only changes at every learning rate schedule. In practice, we noticed accuracy degradation especially when the levels are not updated during the initial iterations where the average variance is rapidly changing.

Convergence of level updating. Fig. 8 shows the expected normalized variance (the objective in (3)) and expected variance (the objective in (9)) during one step of adapting levels. This figure shows that the objective function in (3) is nonconvex and different initializations lead to sub-optimal solutions. ALQG and ALQG-N refer to variations of ALQ using gradient descent instead of coordinate descent.

Hyperparameters used for training. Table 3 shows the hyperparameters used for training CIFAR-10 and ImageNet. These are conventional hyperparameters for training ResNet models. SuperSGD is able to replicate the accuracy reported in [28] showing the correct setting for training.

Validation accuracy on full ImageNet run Table 4 shows the validation accuracy of full ImageNet runs on ResNet-18. Total number of iterations required for a full ImageNet run is 600K. This table shows that ALQ and ALQ-N are able to outperform QSGDinf by 1% on ImageNet.




Figure 8: Convergence of different level update methods

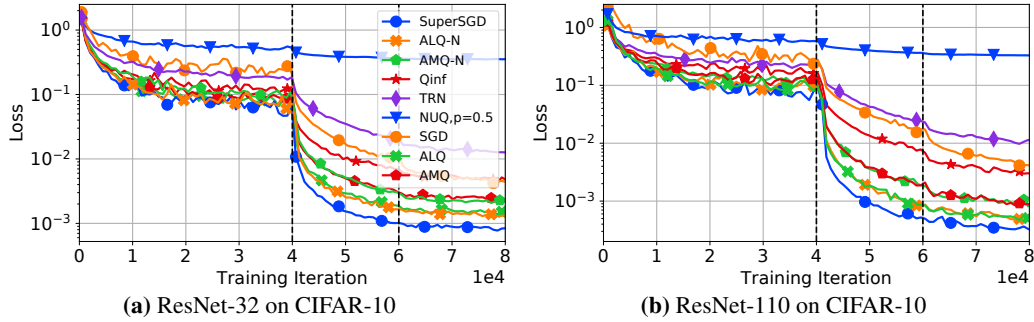


Figure 9: Extended Training loss on CIFAR-10. All methods use 3 quantization bits. Bucket size for ResNet-110 trained on CIFAR-10 is 16384 and for ResNet-32 is 8192.

Similar performance of normalized and unnormalized variations of AMQ and ALQ methods suggests that for given datasets and deep models, the distribution of normalized gradient coordinates can be represented by either of the formulations in Section 3. In AMQ-N and ALQ-N, μ and σ values for the truncated normal distribution is equal to the average of μ and σ for individual buckets.

Figs. 10a and 10b show an interesting observation for NUQSGD. Although NUQSGD has worse performance in terms of the training loss and average variance compared to all other approaches, it is able to achieve better validation accuracy. This suggests that NUQSGD is able to generalize better in this specific setting. However, this pattern does not repeat when it comes to the ImageNet dataset.

K.1 Revised Experiments

Figs. 9 to 12 are extended versions of Figs. 3, 4 and 11 figures in the main body. The difference is that they contain more baselines compared to the figures in the main body. Fig. 9 contains the training loss for the experiments on CIFAR-10. It was not possible to include the same figure for ImageNet, because calculating the full training loss on ImageNet takes a very long time.

The expected variance, training loss, and the validation loss for the results presented in Table 2 are shown in Fig. 13. Although ALQ performs better in expected variance and training loss, it seems to have trouble when it comes to the validation loss for 32-GPUs. We suspect that this is due to the large total batch size used for these experiment that results in overfitting. The batch size for each GPU is 128.

K.2 Effect of Using Gradient Clipping

TRN [15] introduced the idea of gradient clipping before quantization. Gradient clipping replaces the gradient coordinates that are far from the mean to reduce the gradient variance. The gradient coordinates that are very far from the mean can affect the normalization. In order to tackle this




Figure 10: Extended Validation loss on CIFAR-10 and ImageNet. All methods use 3 quantization bits. Bucket size for ResNet-110 trained on CIFAR-10 is 16384, for ResNet-32 is 8192, and for ResNet-18 on ImageNet is 8192.




Figure 11: Extended Variance (no train) on CIFAR-10 and ImageNet. All methods use 3 quantization bits. Bucket size for ResNet-110 trained on CIFAR-10 is 16384, for ResNet-32 is 8192, and for ResNet-18 on ImageNet is 8192.

problem, they clip the gradients before quantization. The clipping process can be described using Eq. (49):

$$f(g_i) = \begin{cases} g_i & |g_i| \leq c\sigma \\ \text{sign}(g_i) \cdot c\sigma & |g_i| > c\sigma \end{cases} \quad (49)$$

The constant c used in TRN equals to 2.5. In order to investigate the effect of gradient clipping in ALQ and AMQ, we train a ResNet-8 on CIFAR-10 dataset for various bucket sizes. Fig. 14 shows the validation accuracy of the base-lines and the algorithms we proposed. ALQ and ALQ-N always maintain better or equal accuracy compared to the other quantization schemes. It is also worth noting that the quantization is performed by each layer instead of a performing the quantization across the network without considering the layers.




Figure 14: Validation Loss

K.3 Timing Overhead

In this section, we provide the timing results per step for training ResNet-18 (Table 6) and ResNet-50 (Table 7) on ImageNet with mini-batch size 512. The training setup consists of 4 AWS nodes with one V100 GPU on each. Network bandwidth programmatically constrained to 1Gbit/s.




Figure 12: Extended Variance on CIFAR-10 and ImageNet. All methods use 3 quantization bits. Bucket size for ResNet-110 trained on CIFAR-10 is 16384, for ResNet-32 is 8192, and for ResNet-18 on ImageNet is 8192.




Figure 13: Using 32-GPUs to train ResNet-32 on CIFAR-10.

Table 5: Training ResNet50 on ImageNet with min-batch size 512. Time per step for training with 32bits full-precision is 1.2s and with 16 bits full-precision is 0.61s.

Bits	Bucket size	Time per step (s)	Ratio to FP32	Ratio to FP16
2	64	0.41	0.34	0.67
2	256	0.39	0.33	0.64
2	1024	0.38	0.32	0.62
2	8192	0.38	0.32	0.62
2	16384	0.38	0.32	0.62
3	64	0.4	0.33	0.66
3	256	0.4	0.33	0.66
3	1024	0.4	0.33	0.66
3	8192	0.4	0.33	0.66
3	16384	0.38	0.32	0.62
4	64	0.42	0.35	0.69
4	256	0.41	0.34	0.67
4	1024	0.4	0.33	0.66
4	8192	0.4	0.33	0.66
4	16384	0.4	0.33	0.66
5	64	0.43	0.36	0.70
5	256	0.42	0.35	0.69
5	1024	0.41	0.34	0.67
5	8192	0.4	0.33	0.66
5	16384	0.4	0.33	0.66
6	64	0.42	0.35	0.69
6	256	0.41	0.34	0.67
6	1024	0.41	0.34	0.67
6	8192	0.41	0.34	0.67
6	16384	0.41	0.34	0.67
7	64	0.45	0.38	0.74
7	256	0.43	0.36	0.70
7	1024	0.42	0.35	0.69
7	8192	0.42	0.35	0.69
7	16384	0.43	0.36	0.70
8	64	0.45	0.38	0.74
8	256	0.44	0.37	0.72
8	1024	0.43	0.36	0.70
8	8192	0.43	0.36	0.70
8	16384	0.43	0.36	0.70

Table 6: Training ResNet18 on ImageNet with min-batch size 512. Time per step for training with 32bits full-precision is 0.57s and with 16 bits full-precision is 0.28s.

Bits	Bucket size	Time per step	Ratio to FP32	Ratio to FP16
2	64	0.13	0.23	0.46
2	256	0.12	0.21	0.43
2	1024	0.11	0.19	0.39
2	8192	0.11	0.19	0.39
2	16384	0.11	0.19	0.39
3	64	0.13	0.23	0.46
3	256	0.12	0.21	0.43
3	1024	0.12	0.21	0.43
3	8192	0.12	0.21	0.43
3	16384	0.12	0.21	0.43
4	64	0.13	0.23	0.46
4	256	0.13	0.23	0.46
4	1024	0.12	0.21	0.43
4	8192	0.12	0.21	0.43
4	16384	0.12	0.21	0.43
5	64	0.13	0.23	0.46
5	256	0.13	0.23	0.46
5	1024	0.13	0.23	0.46
5	8192	0.13	0.23	0.46
5	16384	0.13	0.23	0.46
6	64	0.14	0.25	0.50
6	256	0.13	0.23	0.46
6	1024	0.13	0.23	0.46
6	8192	0.13	0.23	0.46
6	16384	0.13	0.23	0.46
7	64	0.14	0.25	0.50
7	256	0.13	0.23	0.46
7	1024	0.14	0.25	0.50
7	8192	0.13	0.23	0.46
7	16384	0.13	0.23	0.46
8	64	0.15	0.26	0.54
8	256	0.14	0.25	0.50
8	1024	0.14	0.25	0.50
8	8192	0.14	0.25	0.50
8	16384	0.14	0.25	0.50

Table 7: Additional overhead of proposed methods for training ResNet18 on ImageNet (Table 6). We also show the cost of performing 3 updates relative to the total cost of training for 60 epochs. Full-precision training for 60 epochs with 32 bits takes 95 hours while with 16 bits takes 46 hours.

Bits	Bucket size	Quantization Method	Time per update	Ratio to FP32	Ratio to FP16
3	64	ALQ-N	1012	0.89	1.81
3	256	ALQ-N	630	0.55	1.13
3	1024	ALQ-N	533	0.47	0.95
3	8192	ALQ-N	559	0.49	1.00
3	16384	ALQ-N	591	0.52	1.06
4	64	ALQ-N	1170	1.03	2.09
4	256	ALQ-N	822	0.72	1.47
4	1024	ALQ-N	733	0.64	1.31
4	8192	ALQ-N	681	0.60	1.22
4	16384	ALQ-N	684	0.60	1.22
6	64	ALQ-N	2036	1.79	3.64
6	256	ALQ-N	1710	1.50	3.05
6	1024	ALQ-N	1556	1.36	2.78
6	8192	ALQ-N	1574	1.38	2.81
6	16384	ALQ-N	1671	1.47	2.98
8	64	ALQ-N	5604	4.92	10.01
8	256	ALQ-N	5253	4.61	9.38
8	1024	ALQ-N	5478	4.81	9.78
8	8192	ALQ-N	5180	4.54	9.25
8	16384	ALQ-N	5576	4.89	9.96
3	64	ALQ	1032	0.91	1.84
3	256	ALQ	585	0.51	1.04
3	1024	ALQ	444	0.39	0.79
3	8192	ALQ	474	0.42	0.85
3	16384	ALQ	477	0.42	0.85
4	64	ALQ	930	0.82	1.66
4	256	ALQ	529	0.46	0.94
4	1024	ALQ	450	0.39	0.80
4	8192	ALQ	431	0.38	0.77
4	16384	ALQ	486	0.43	0.87
6	64	ALQ	974	0.85	1.74
6	256	ALQ	573	0.50	1.02
6	1024	ALQ	489	0.43	0.87
6	8192	ALQ	428	0.38	0.76
6	16384	ALQ	438	0.38	0.78
8	64	ALQ	1051	0.92	1.88
8	256	ALQ	637	0.56	1.14
8	1024	ALQ	516	0.45	0.92
8	8192	ALQ	508	0.45	0.91
8	16384	ALQ	516	0.45	0.92