

# Hardware-Efficient Mixed-Precision CP Tensor Decomposition \*

Zi Yang, Junnan Shan, and Zheng Zhang

**Abstract.** Tensor decomposition has been widely used in machine learning and high-volume data analysis. However, large-scale tensor factorization often consumes huge memory and computing cost. Meanwhile, modernized computing hardware such as tensor processing units (TPU) and Tensor Core GPU has opened a new window of hardware-efficient computing via mixed- or low-precision arithmetic representations. In this paper, we exploit the low-precision representation of tensor factorization, and propose a mixed-precision block stochastic gradient descent (SGD) method to reduce the costs of CP tensor decomposition. Our method achieves robust and fast convergence via a two-stage optimization, i.e., SignSGD followed by mixed-precision SGD. Detailed theoretical analysis is provided to prove the convergence of the proposed mixed-precision algorithm. Numerical experiments on both synthetic and realistic tensor data sets show the superior efficiency of our mixed-precision algorithm compared to full-precision CP decomposition. This work can remarkably reduce the memory, computing and energy cost on resource-constraint edge computing devices. We demonstrate this benefit via an FPGA prototype.

**Key words.** Tensor decomposition, mixed-precision optimization, stochastic gradients, hardware-aware algorithms.

**MSC codes.** 15A69, 90C15

**1. Introduction.** As a higher-order generalization of matrices, tensors [35] have been used to represent and process multi-dimensional arrays in many science and engineering fields, including quantum physics [20, 30, 40, 45], scientific computing [6, 47], uncertainty quantification [18, 55, 56], machine learning [2, 23, 25, 26, 41, 42, 33, 49] and many others. Many successful applications rely on efficient tensor decompositions [8, 15, 22, 46], which represent an original high-order high-volume data array with some low-rank factors to achieve huge memory and computing cost reduction. For instance, tensor decomposition has achieved orders-of-magnitude parameter reduction of deep neural networks [25, 26, 33, 42], enabling their energy-efficient training and deployment on edge devices. As one of the most popular tensor decomposition methods, the CANDECOMP/PARAFAC (CP) decomposition [8] factorizes a large tensor into the summation of some rank-1 tensors. A CP factorization is often obtained via algebraic methods [19, 39] or numerical optimization techniques such as gradient-based optimization [21] and alternating minimization [13]. The former provides excellent theoretical guarantees, but are neither noise-resistant nor scalable to high tensor ranks. The later has better efficiency, but computing the full gradients is expensive for high-volume tensor data sets. Motivated by the success in large-scale machine learning, recent approaches use stochastic gradient descent (SGD) methods [3, 5, 36, 51] to relief the high computation cost in tensor factorization. So far, most (if not all) tensor decomposition algorithms are developed for classical computing platforms (e.g., CPU and conventional GPU) that use double-precision 64-bit or single-precision 32-bit floating-point data representations.

---

\*The authors are with Department of Electrical and Computer Engineering, University of California at Santa Barbara, CA. (Emails: ziy@ucsb.edu, junnanshan@ucsb.edu, zzhang01@ucsb.edu).

On the other hand, the recent revolution of artificial intelligence has triggered massive interests in computing hardware that supports mixed-precision and low-precision computation. For instance, Google’s Tensor Processing Units (TPUs) [31] can easily handle machine learning tasks with 16-bit floating point representations. NVIDIA’s tensor Core GPU supports double-, single- and half-precision floating-point operations, as well as various low-precision integer operations. Reconfigurable computing platforms such as field-programmable gate arrays (FPGA) can support arbitrarily low-precision computation to save energy and hardware utilization. These mixed-precision computing platforms are very suitable for the training and inference of deep learning models [17, 16, 29, 50, 54], due to their error-resilient activation functions or output operators. Interestingly, recently mixed-precision computing has also shown great success in many scientific computing tasks [1, 9, 10, 11, 12, 24, 44] such as LU factorization, Cholesky factorization, least square optimization, GMRES. However, mixed-precision computing has been rarely investigated for tensor computation. We envision that similar memory and runtime benefit can be obtained by developing mixed-precision tensor computation algorithms. As the development of 5G and future 6G networks, more and more (possibly private and sensitive) data needs to be processed on resource-constraint edge devices, where mixed-precision tensor computation will play an increasingly important role.

In this paper, we make the first step of exploring low-precision tensor computation by proposing a novel mixed-precision CP tensor decomposition algorithm. By utilizing low-precision stochastic gradient computation in a two-stage optimization framework, our method can remarkably reduce the computation and energy costs of CP decomposition. Our main contributions are summarized below.

- We propose a computationally efficient CP decomposition via a mixed-precision SGD method. We improve the convergence via a mixed-precision SignSGD initialization. We carefully design the low-precision stochastic gradient computation to maximize the computational efficiency and minimize the accuracy drop via analyzing the sensitivity of each step with respect to the quantization errors.
- We prove the convergence of the proposed mixed-precision CP decomposition. Under some conditions, we firstly show that the CP decomposition problem is locally strongly convex after proper normalization. Then, we prove that SignSGD with mixed-precision gradients converges to a stationary point up to a noise level. Finally, we prove that the mixed-precision SGD has a locally linear convergence rate for our problem.
- Numerical experiments demonstrate that our mixed-precision approach can remarkably reduce the computation cost while attaining similar accuracy to the full-precision algorithm. An FPGA prototype further demonstrates the saving of run-time, hardware resources, and energy on edge computing devices.

We remark that the proposed mixed-precision stochastic gradient can be applied to all SGD-based algorithms for CP tensor decomposition.

## 2. Preliminary.

**Notation.** Throughout the paper, lower-case letters (e.g.,  $a$ ) denote scalars; lower-case bold letters (e.g.,  $\mathbf{a}$ ) denote vectors; upper case bold letters (e.g.,  $\mathbf{A}$ ) denote matrices. We use  $\mathbf{1}$  or  $\mathbf{0}$  to denote a vector/matrix whose entries are all 1 or 0, respectively.  $\mathbf{I}_n$  is an  $n$ -by- $n$  identity matrix. We use upper-case calligraphic bold letters (e.g.,  $\mathcal{A}$ ) to denote tensors,

which are high-dimensional generalizations of matrices. We use  $[n]$  to denote the set of integers  $\{1, 2, \dots, n\}$ . For a vector  $\mathbf{v}$ ,  $\|\mathbf{v}\|$  and  $\|\mathbf{v}\|_1$  denote its Euclidean norm and 1-norm, respectively. For a matrix  $\mathbf{A}$ ,  $\mathbf{A}^T$  denotes its transpose;  $\text{tr}(\mathbf{A})$  denotes the trace of  $\mathbf{A}$ ;  $\|\mathbf{A}\|$  represents the Frobenius norm, and the spectrum norm  $\|\mathbf{A}\|_2$  is the largest singular value of  $\mathbf{A}$ . We use MATLAB-style indexing to denote submatrices. For instance,  $\mathbf{A}(i_1 : i_2, j_1 : j_2)$  denotes the submatrix consisting of the rows from  $i_1$  to  $i_2$  and the columns from  $j_1$  to  $j_2$ . The function  $\text{sign}(a)$  obtains the sign of  $a$ , i.e.,

$$\text{sign}(a) := \begin{cases} 1 & \text{if } a > 0 \\ 0 & \text{if } a = 0 \\ -1 & \text{if } a < 0 \end{cases}.$$

The sign function can be used for matrices and tensors by applying the function element-wisely.

For a twice-differentiable function  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ , we use  $\nabla f \in \mathbb{R}^n$  and  $\nabla^2 f \in \mathbb{R}^{n \times n}$  to denote the gradient and the Hessian matrix of  $f$ , respectively. The function  $f$  is  $\lambda$ -strongly convex for  $\lambda > 0$  if the smallest eigenvalue of  $\nabla^2 f$  is not less than  $\lambda$ . Equivalently,  $f$  is  $\lambda$ -strongly convex if

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \nabla f(\mathbf{x})^T(\mathbf{y} - \mathbf{x}) + \frac{\lambda}{2}\|\mathbf{y} - \mathbf{x}\|^2, \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n.$$

For the vector valued function  $\mathbf{h}(\mathbf{x}) = (h_1(\mathbf{x}), \dots, h_m(\mathbf{x}))$  where  $h_i : \mathbb{R}^n \rightarrow \mathbb{R}$ , the Jacobian matrix  $\mathbf{J}_{\mathbf{h}}(\mathbf{x})$  is

$$\mathbf{J}_{\mathbf{h}}(\mathbf{x}) := (\nabla h_1(\mathbf{x}), \dots, \nabla h_m(\mathbf{x}))^T.$$

**2.1. Tensors.** Tensors can be regarded as multi-dimensional data arrays [37]. The space of real tensors with order  $m$  and dimension  $N_1, N_2, \dots, N_m$  is denoted by  $\mathbb{R}^{N_1 \times N_2 \times \dots \times N_m}$ . The  $(i_1, i_2, \dots, i_m)$ -th element of a tensor  $\mathcal{A} \in \mathbb{R}^{N_1 \times N_2 \times \dots \times N_m}$  is denoted as  $a_{i_1, \dots, i_m}$  for  $1 \leq i_j \leq N_j$ . The mode- $k$  unfolding of  $\mathcal{A}$  is the matrix  $\mathbf{A}_{[k]} \in \mathbb{R}^{N_k \times (\prod_{i=1}^m N_i)/N_k}$  which is obtained by reshaping  $\mathcal{A}$  with the  $k$ th dimension being the leading dimension. The Frobenius norm of  $\mathcal{A}$  is

$$\|\mathcal{A}\|_{\text{F}} := \sqrt{\sum_{i_1, \dots, i_m} a_{i_1, \dots, i_m}^2}.$$

For vectors  $\{\mathbf{u}_i \in \mathbb{R}^{N_i}\}_{i=1}^m$ , their outer product forms an order- $m$  rank-1 tensor

$$\mathcal{B} = \mathbf{u}_1 \circ \mathbf{u}_2 \circ \dots \circ \mathbf{u}_m \iff b_{i_1, \dots, i_m} = \prod_{k=1}^m \mathbf{u}_k(i_k).$$

A tensor  $\mathcal{A} \in \mathbb{R}^{N_1 \times N_2 \times \dots \times N_m}$  is said to have a rank- $r$  CP decomposition if there exist matrices  $\{\mathbf{U}_i \in \mathbb{R}^{N_i \times r}\}_{i=1}^m$  such that

$$\mathcal{A} = [\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_m] := \sum_{j=1}^r \mathbf{U}_1(:, j) \circ \dots \circ \mathbf{U}_m(:, j).$$

Type	Bits	Sign	Exponent	Significand	Min	Max
FP <sub>16</sub>	16	1	5	10	$6.1 \times 10^{-5}$	$6.6 \times 10^4$
FP <sub>32</sub>	32	1	8	23	$1.2 \times 10^{-38}$	$3.4 \times 10^{38}$
FP <sub>64</sub>	64	1	11	52	$2.2 \times 10^{-308}$	$1.8 \times 10^{308}$

**Table 1**  
Floating Point Representations

$(m, k, n)$	INT <sub>8</sub> time	FP <sub>16</sub> time	FP <sub>32</sub> time
(240, 240 <sup>2</sup> , 256)	232 ( <b>4.47</b> ×)	675 (1.54×)	1037 (1×)
(60, 60 <sup>3</sup> , 64)	794 ( <b>4.55</b> ×)	2456 (1.47×)	3615 (1×)
(24, 24 <sup>4</sup> , 32)	1139 ( <b>4.89</b> ×)	3784 (1.47×)	5571 (1×)

**Table 2**

Time comparisons of matrix multiplications of  $m \times k$  and  $k \times n$  under various precisions on GPU. The times are measured in microseconds ( $\mu$ s).

The smallest integer  $r$  that ensures the above equality is called the CP rank of  $\mathcal{A}$ , denoted by  $\text{rank}(\mathcal{A})$ . The Khatri-Rao product of matrices  $\mathbf{U}_1, \dots, \mathbf{U}_m$  is a column-wise Kronecker product, i.e.,

$$\mathbf{U}_1 \odot \dots \odot \mathbf{U}_m := [\otimes_{i=1}^m \mathbf{U}_i(:, 1), \dots, \otimes_{i=1}^m \mathbf{U}_i(:, r)],$$

where  $\otimes$  denotes the Kronecker product. It holds that  $\mathbf{A}_{[k]} = \mathbf{U}_k(\odot_{i=1, i \neq k}^m \mathbf{U}_i)^T$ .

**2.2. Precision Representations.** In practice, numbers are represented and processed as binary strings on digital computing hardware. The binary strings can represent numbers in either fixed-point format or floating-point format. We use INT <sub>$n$</sub>  and FP <sub>$n$</sub>  to denote an  $n$ -bit fixed-point format and an  $n$ -bit floating-point format, respectively. The representation format is directly related to the precision of the represented number. Hence, the representation format is also called precision format.

An INT <sub>$n$</sub>  data representation uses  $n$  bits, where the first bit stores the sign and the other  $n - 1$  bits store the absolute value. The set of numbers that the INT <sub>$n$</sub>  format can represent is  $\{-2^{n-1}, -2^{n-1} + 1, \dots, 0, 1, \dots, 2^{n-1} - 1\}$ . The FP <sub>$n$</sub>  format uses 1 bit to store the sign,  $N_1$  bits to store significand, and  $N_2$  bits to store exponent, where  $N_1 + N_2 + 1 = n$ . Then, the number is represented by  $\text{sign} \times \text{significand} \times 2^{\text{exponent}}$ . FP<sub>16</sub> (half precision), FP<sub>32</sub> (single precision), and FP<sub>64</sub> (double precision) are most commonly used and are supported by most devices. Their bits for each part and representation ranges are described in Table 1. Floating-point arithmetic operations are much more expensive than fixed-point arithmetic operations with the same number of bits. Clearly, low-bit representations consume less memory and computation resources but cause larger rounding-off errors. Table 2 compares the run-time of matrix multiplications under different precision formats on tensor core GPU. The chosen test sizes are common in computing gradients for the proposed Algorithm 3.1 as in (3.5). We can see that the INT<sub>8</sub> multiplications are 4× to 5× faster than FP<sub>32</sub> multiplications.

Deterministic rounding and stochastic rounding methods can be used to round a high-precision number to a lower precision. For a given precision format  $p$ , let  $\mathcal{R}(p)$  be the set of numbers that can be represented by the format  $p$ . The ceiling and floor functions with

precision  $p$  are defined as

$$\begin{aligned} \lceil y \rceil_p &:= \min\{v \in \mathcal{R}(p) \cup \{+\infty\} | v \geq y\}, \\ \lfloor y \rfloor_p &:= \max\{v \in \mathcal{R}(p) \cup \{-\infty\} | v \leq y\}. \end{aligned}$$

When the precision  $p$  is not specified, we use  $\mathcal{R}(p) = \mathbb{N}$  by default. The quantization function  $\mathbf{Q}_{p,\delta}^D$ , with precision  $p$ , scaling factor  $\delta$ , and deterministic rounding, is defined as

$$\mathbf{Q}_{p,\delta}^D(x) = \begin{cases} \delta \lceil x/\delta \rceil_p & \text{if } x/\delta \geq \frac{1}{2} (\lceil x/\delta \rceil_p + \lfloor x/\delta \rfloor_p) \\ \delta \lfloor x/\delta \rfloor_p & \text{if } x/\delta < \frac{1}{2} (\lceil x/\delta \rceil_p + \lfloor x/\delta \rfloor_p). \end{cases}$$

The quantization function  $\mathbf{Q}_{p,\delta}^S$  with stochastic rounding is

$$\mathbf{Q}_{p,\delta}^S(x) = \begin{cases} \delta \lceil x/\delta \rceil_p & \text{with probability } \frac{x/\delta - \lfloor x/\delta \rfloor_p}{\lceil x/\delta \rceil_p - \lfloor x/\delta \rfloor_p} \\ \delta \lfloor x/\delta \rfloor_p & \text{with probability } \frac{\lceil x/\delta \rceil_p - x/\delta}{\lceil x/\delta \rceil_p - \lfloor x/\delta \rfloor_p} \end{cases}.$$

The stochastic rounding ensures that the quantization is unbiased, i.e.,  $\mathbb{E}(\mathbf{Q}_{p,\delta}^S(x)|x) = x$ .

**3. Proposed Algorithm.** This section presents a mixed-precision SGD-type algorithm to reduce the memory and computation cost of CP tensor decomposition. This method has a linear convergence rate when it gets close to the optimal solution. A mixed-precision SignSGD method is utilized at the beginning to improve the convergence of the whole framework.

**3.1. Mixed-Precision CP Decomposition.** Given a tensor  $\mathcal{A} \in \mathbb{R}^{N_1 \times \dots \times N_m}$ , the rank- $r$  CP tensor decomposition can be formulated as the optimization problem

$$(3.1) \quad \min_{\Theta} f(\Theta) = \|\mathcal{A} - \llbracket \mathbf{U}_1, \dots, \mathbf{U}_m \rrbracket\|_F^2, \text{ with } \Theta = \{\mathbf{U}_i \in \mathbb{R}^{N_i \times r}\}_{i=1}^m.$$

This problem can be rewritten as

$$(3.2) \quad \min_{\Theta} f := \frac{1}{N} \sum_{i_1=1}^{N_1} \dots \sum_{i_m=1}^{N_m} (a_{i_1 \dots i_m} - \llbracket \mathbf{U}_1(i_1, :), \dots, \mathbf{U}_m(i_m, :) \rrbracket)^2,$$

where  $N := N_1 \dots N_m$ . Since the cost function is the summation of  $N$  functions, we can naturally apply an SGD-type method to solve the optimization.

Instead of using standard SGD [7], we present a mixed-precision SGD-type algorithm to solve Problem (3.1). Let  $\mathbf{U}_1^s, \dots, \mathbf{U}_m^s$  be the tensor factor matrices in the  $s$ -th iteration and  $\mathbf{Q}(\tilde{\mathbf{g}}_i^s)$  be the quantized stochastic gradient with respect to  $\mathbf{U}_i^s$ . Corollary 4.2 shows that Problem (3.1) is locally strongly convex around the true decomposition if the leading rows of  $\mathbf{U}_i(1, :)$  are fixed for  $i = 2, \dots, m$ . We propose to update variables as

$$\mathbf{U}_1^{s+1} = \mathbf{U}_1^s - \alpha_s \mathbf{Q}(\tilde{\mathbf{g}}_1^s),$$

$$\mathbf{U}_i^{s+1}(2 : N_i, :) = \mathbf{U}_i^s(2 : N_i, :) - \alpha_s \mathbf{Q}(\tilde{\mathbf{g}}_i^s)(2 : N_i, :).$$

---

**Algorithm 3.1** Mixed-Precision Stochastic Gradient Algorithm for Tensor Decomposition

---

- 1: **Input:** tensor  $\mathcal{A} \in \mathbb{R}^{N_1 \times \dots \times N_m}$ , rank  $r$ , initialization  $\{\mathbf{U}_i^0 \in \mathbb{R}^{N_i \times r}\}_{i=1}^m$ , initial learning rates  $\alpha_0^{\text{sign}}, \alpha_0^{\text{SGD}} > 0$ , thresholds  $\epsilon_1 > \epsilon_2 > 0$ , quantizations  $\mathbf{Q}_1, \mathbf{Q}_2$ , sample sizes  $\{n_i\}_{i=1}^m$ , learning rate update intervals  $K^{\text{sign}}, K^{\text{SGD}}$ , learning rate update constants  $\eta^{\text{sign}}, \eta^{\text{SGD}}$ .
  - 2: Let  $s = 0$ .
  - 3:  $\alpha_0 = \alpha_0^{\text{sign}}$ .
  - 4: **while**  $\|\mathcal{A} - \llbracket \mathbf{U}_1^s, \dots, \mathbf{U}_m^s \rrbracket\| / \|\mathcal{A}\| > \epsilon_1$  **do**
  - 5:   Compute the mixed-precision gradient  $\{\mathbf{Q}(\tilde{\mathbf{g}}_i^s)\}_{i=1}^m$  as in Algorithm 3.2.
  - 6:    $\mathbf{U}_i^{s+1} = \mathbf{U}_i^s - \alpha_s \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s))$ .
  - 7:    $\alpha_{s+1} = \begin{cases} \eta^{\text{sign}} \alpha_s, & \text{if } ((s+1) \bmod K^{\text{sign}}) = 0 \\ \alpha_s, & \text{otherwise} \end{cases}$ .
  - 8:    $s = s + 1$ .
  - 9: **end while**
  - 10: Let  $\alpha_s = \alpha_0^{\text{SGD}}$ ,  $s^{\text{sign}} = s$ .
  - 11: **while**  $\|\mathcal{A} - \llbracket \mathbf{U}_1^s, \dots, \mathbf{U}_m^s \rrbracket\| / \|\mathcal{A}\| > \epsilon_2$ . **do**
  - 12:   Compute the mixed-precision gradient  $\{\mathbf{Q}(\tilde{\mathbf{g}}_i^s)\}_{i=1}^m$  as in Algorithm 3.2.
  - 13:    $\mathbf{U}_1^{s+1} = \mathbf{U}_1^s - \alpha_s \mathbf{Q}(\tilde{\mathbf{g}}_1^s)$ .
  - 14:    $\mathbf{U}_i^{s+1}(2:N_i, :) = \mathbf{U}_i^s(2:N_i, :) - \alpha_s \mathbf{Q}(\tilde{\mathbf{g}}_i^s)(2:N_i, :)$ ,  $i = 2, \dots, m$
  - 15:    $\alpha_{s+1} = \begin{cases} \eta^{\text{SGD}} \alpha_s, & \text{if } ((s - s^{\text{sign}} + 1) \bmod K^{\text{SGD}}) = 0 \\ \alpha_s, & \text{otherwise} \end{cases}$ .
  - 16:    $s = s + 1$ .
  - 17: **end while**
  - 18: **Output:** factor matrices  $\{\mathbf{U}_i^s\}_{i=1}^m$ .
- 

Problem (3.1) has many stationary points, and the mixed-precision SGD can easily converge to a local optimizer without a good initialization point. We propose to use mixed-precision SignSGD to find a good initialization for SGD, which updates variables as follows:

$$\mathbf{U}_i^{s+1} = \mathbf{U}_i^s - \alpha_s \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s)).$$

The mixed-precision SignSGD only uses the sign of the gradient to update parameters. Consequently, it is more robust against non-convexity and quantization errors. In practice, we find that SignSGD is unlikely to be trapped by a stationary point. This motivates us to firstly run mixed-precision SignSGD for a number of iterations. When the error becomes small, we switch to mixed-precision SGD for better accuracy and faster convergence.

The learning rate is updated as  $\alpha_{s+1} = \eta \alpha_s$  every  $K$  iterations for some constant  $1 > \eta > 0$ . It is the multi-stage update rule. The complete mixed-precision CP decomposition (3.1) is presented in Algorithm 3.1.

**3.2. Mixed-Precision Block Stochastic Gradient.** Gradient computation is often the most expensive part in SGD-type algorithms. This subsection describes how to efficiently compute the mixed-precision stochastic gradient used in Algorithm 3.1.

Problem (3.2) is well-structured, therefore we use block sampling to maximize the usage of parallel computing. In each iteration, we uniformly sample a subset of indices  $\mathcal{I}_i \subset [N_i]$

---

**Algorithm 3.2** Compute Mixed-Precision Stochastic Gradient
 

---

- 1: **Input:** tensor  $\mathcal{A} \in \mathbb{R}^{N_1 \times \dots \times N_m}$ , rank  $r$ , factor matrices  $\{\mathbf{U}_i^0 \in \mathbb{R}^{N_i \times r}\}_{i=1}^m$ , quantization functions  $\mathbf{Q}_1, \mathbf{Q}_2$ , sample sizes  $\{n_i\}_{i=1}^m$ .
  - 2: Randomly sample  $\mathcal{I}_i \subset [n_i]$  with  $|\mathcal{I}_i| = n_i$  for  $i \in [m]$ .
  - 3: Compute  $\mathcal{M} = -\mathcal{A}(\mathcal{I}_1, \dots, \mathcal{I}_m) + \llbracket \mathbf{Q}_1(\mathbf{U}_1(\mathcal{I}_1, :)), \dots, \mathbf{Q}_1(\mathbf{U}_m(\mathcal{I}_m, :)) \rrbracket$ .
  - 4: Compute  $\mathbf{V}_i = \odot_{j=1, j \neq i}^m \mathbf{Q}_1(\mathbf{U}_j(\mathcal{I}_j, :))$  for  $i \in [m]$ .
  - 5: Compute  $\mathbf{Q}(\tilde{\mathbf{g}}_i)(\mathcal{I}_i, :) = \mathbf{Q}_2(\mathbf{M}_{[i]})\mathbf{Q}_2(\mathbf{V}_i)$ .
  - 6: **Output:** Mixed-precision stochastic gradient  $\{\mathbf{Q}(\tilde{\mathbf{g}}_i)\}_{i=1}^m$ .
- 

with  $|\mathcal{I}_i| = n_i$  for  $i = 1, \dots, m$ . Then, we consider the cost function

$$f_{\mathcal{I}} := \frac{1}{n} \|\mathcal{A}(\mathcal{I}_1, \dots, \mathcal{I}_m) - \llbracket \mathbf{U}_1(\mathcal{I}_1, :), \dots, \mathbf{U}_m(\mathcal{I}_m, :) \rrbracket\|_{\mathbb{F}}^2, \text{ with } n = n_1 \cdots n_m.$$

The gradient of  $f_{\mathcal{I}}$  with respect to  $\mathbf{U}_i(\mathcal{I}_i, :)$  is

$$(3.3) \quad \nabla_{\mathbf{U}_i(\mathcal{I}_i, :)} f_{\mathcal{I}} = -\frac{2}{n} (\mathcal{A}(\mathcal{I}_1, \dots, \mathcal{I}_m) - \llbracket \mathbf{U}_1(\mathcal{I}_1, :), \dots, \mathbf{U}_m(\mathcal{I}_m, :) \rrbracket)_{[i]} \odot_{j=1, j \neq i}^m \mathbf{U}_j(\mathcal{I}_j, :).$$

Therefore, the stochastic gradient  $\tilde{\mathbf{g}}_i := \nabla_{\mathbf{U}_i} f_{\mathcal{I}} \in \mathbb{R}^{N_i \times r}$  is given as

$$(3.4) \quad \tilde{\mathbf{g}}_i(j_i, :) = \begin{cases} \nabla_{\mathbf{U}_i(j_i, :)} f_{\mathcal{I}} & \text{if } j_i \in \mathcal{I}_i \\ 0 & \text{if } j_i \notin \mathcal{I}_i \end{cases}.$$

We regard  $\mathcal{I} := (\mathcal{I}_1, \dots, \mathcal{I}_m)$  as a random variable. Each  $\mathcal{I}_i$  is sampled uniformly, hence it holds that  $\mathbb{E}_{\mathcal{I}}[\tilde{\mathbf{g}}_i] = \mathbf{g}_i$  for  $i = 1, \dots, m$ , where  $\mathbf{g}_i := \nabla_{\mathbf{U}_i} f$ .

We compute the quantized value of the block stochastic gradient  $\tilde{\mathbf{g}}_i$  (3.4) as follows:

$$(3.5a) \quad \mathcal{M} := -\mathcal{A}(\mathcal{I}_1, \dots, \mathcal{I}_m) + \llbracket \mathbf{Q}_1(\mathbf{U}_1(\mathcal{I}_1, :)), \dots, \mathbf{Q}_1(\mathbf{U}_m(\mathcal{I}_m, :)) \rrbracket,$$

$$(3.5b) \quad \mathbf{V}_i := \odot_{j=1, j \neq i}^m \mathbf{Q}_1(\mathbf{U}_j(\mathcal{I}_j, :)),$$

$$(3.5c) \quad \mathbf{Q}(\tilde{\mathbf{g}}_i)(\mathcal{I}_i, :) := \mathbf{Q}_2(\mathbf{M}_{[i]})\mathbf{Q}_2(\mathbf{V}_i),$$

where  $\mathbf{Q}_1, \mathbf{Q}_2$  are two quantization functions as described in Section 2.2. The steps for computing the stochastic gradient in mixed-precision are summarized in Algorithm 3.2.

The subtraction in (3.5a) and the Khatri–Rao product in (3.5b) are both sensitive to quantization errors, and extremely low-precision quantization function  $\mathbf{Q}_1$  will cause bad convergence behavior. Therefore, we use precision  $\text{FP}_{16}$  and scale  $\delta = 1$  for  $\mathbf{Q}_1$ , i.e.,  $\mathbf{Q}_1 = \mathbf{Q}_{\text{FP}_{16}, 1}$ . The last matrix multiplication (3.5c) is more robust against errors. Consequently, the quantization functions  $\mathbf{Q}_2$  can use an extremely low precision. Practically,  $\text{INT}_4$  and  $\text{INT}_8$  always work well, and  $\text{INT}_2$  can work when the tensor rank  $r$  is small. For the specific quantization  $\mathbf{Q}_{\text{INT}_b, \delta}(\mathbf{X})$  for a matrix  $\mathbf{X}$ , the scaling factor  $\delta$  depends on  $\mathbf{X}$  and the precision  $\text{INT}_b$ . We typically set  $\delta$  slightly less than  $\frac{\max\{\mathbf{X}\}}{2^{b-1}-1}$ . This ensures most entries of  $\mathbf{X}$  lie in the representation range of  $\text{INT}_b$  while preserving low quantization errors.

	INT <sub>8</sub>	INT <sub>4</sub>	INT <sub>2</sub>
$m = 3$	21.9%	17.2%	14.8%
$m = 4$	20.0%	15.0%	12.5%
$m = 5$	18.75%	13.5%	10.9%

**Table 3**

Normalized computation cost compared with full-precision for various orders and precisions.

**Complexity Analysis.** The sub-tensor  $\mathcal{M}$  in (3.5a) is only computed once for all  $i \in [m]$ , and the computation requires around  $2nr$  arithmetic operations. Computing each  $\mathbf{V}_i$  in (3.5b) needs  $\frac{N}{n_i}r$  arithmetic operations, so the total number of operations of step (3.5b) is  $\sum_{i=1}^m \frac{n}{n_i}r$ . Step (3.5c) involves a tensor unfolding along its  $i$ th dimension. The matrix multiplication (3.5c) for each  $i$  requires about  $2nr$  operations. In total, we will do  $m$  such multiplications and the total number of operations is  $2mnr$ . Therefore, the most expensive step in (3.5) is the matrix multiplications (3.5c). Fortunately, (3.5c) is robust against quantization noises, and its cost can be reduced significantly by using ultra low-precision quantization functions. Suppose that each arithmetic operation of precision  $p$  costs  $c_p$  computation resources. Computing the mixed-precision gradient as in (3.5) requires  $C(p_1, p_2) = c_{p_1} \left( 2nr + \sum_{i=1}^m \frac{n}{n_i}r \right) + 2c_{p_2}mnr$  resources, where  $p_1, p_2$  are the precision formats used by  $\mathbf{Q}_1, \mathbf{Q}_2$  respectively. In practice, we typically choose  $p_1$  as FP<sub>16</sub> and  $p_2$  as some low-bit fixed-point format. The computation resource consumed by a specific representation format is proportional to the number of bits. On modern hardware, fixed-point operations typically use less resources and are much faster than floating-point operations. More specifically, fixed-point operations use less than half resources of floating-point operations with the same number of bits [27]. Therefore, we have the estimation  $c_{\text{FP}_{16}} \approx \frac{1}{2}c_{\text{FP}_{32}}$ ,  $c_{\text{INT}_b} \approx \frac{b}{64}c_{\text{FP}_{32}}$ . Then, the estimated costs of (3.5) under full-precision and low-precision are

$$C(\text{FP}_{32}, \text{FP}_{32}) \approx (2 + 2m + \sum_{i=1}^m \frac{1}{n_i})nrc_{\text{FP}_{32}} \approx (2 + 2m)nrc_{\text{FP}_{32}},$$

$$C(\text{FP}_{16}, \text{INT}_b) \approx (1 + \frac{b}{32}m + \sum_{i=1}^m \frac{1}{n_i})nrc_{\text{FP}_{32}} \approx (1 + \frac{b}{32}m)nrc_{\text{FP}_{32}}.$$

The computation saving of using mixed-precision is

$$C(\text{FP}_{16}, \text{INT}_b)/C(\text{FP}_{32}, \text{FP}_{32}) \approx \frac{1 + \frac{b}{32}m}{2 + 2m}.$$

The cost reduction of our proposed mixed-precision gradient is more obvious for high-order tensors and smaller number of bits. Table 3 shows the normalized computational cost for orders  $m = 3, 4, 5$  and precision INT<sub>8</sub>, INT<sub>4</sub>, INT<sub>2</sub>, respectively.

**4. Convergence Analysis.** This section presents the convergence result of Algorithm 3.1. Under some generic conditions, we prove that the tensor decomposition problem (3.1) is locally strongly convex after proper normalization. We prove that the mixed-precision SignSGD



converges to some stationary points up to some noise caused by stochasticity and quantization errors. We also prove that the mixed-precision SGD has a locally linear convergence rate around the global minimizer.

**4.1. Locally Strong Convexity.** This subsection shows the locally strong convexity of the problem (3.1) after proper normalization. Note that Problem (3.1) itself is non-convex and it does not have local convexity as well. Suppose that the tensor  $\mathcal{A}$  has the CP decomposition  $\mathcal{A} = \llbracket \mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_m \rrbracket$ . Then, it holds that

$$(4.1) \quad \mathcal{A} = \sum_{j=1}^r c_{1,j} \mathbf{U}_1(:, j) \circ \dots \circ c_{m,j} \mathbf{U}_m(:, j),$$

for any  $c_{i,j}$ 's as long as  $\prod_{i=1}^m c_{i,j} = 1$ . Therefore, the CP decomposition problem (3.1) has an infinite number of minimizers, but many solutions differ only with scaling factors. Therefore, we fix the elements  $\{\mathbf{U}_i(1, :)\}_{i=2}^m$  and assume that  $\mathbf{U}_i(1, :) = \mathbf{1}^T$  for  $2 \leq i \leq m$  without loss of generality. The CP decomposition problem (3.1) now becomes

$$(4.2) \quad \min_{\tilde{\Theta}} \tilde{f}(\tilde{\Theta}) := \left\| \mathcal{A} - \llbracket \mathbf{U}_1, [\mathbf{1}^T; \tilde{\mathbf{U}}_2], \dots, [\mathbf{1}^T; \tilde{\mathbf{U}}_m] \rrbracket \right\|^2,$$

where  $\tilde{\Theta} := (\mathbf{U}_1, \tilde{\mathbf{U}}_2, \dots, \tilde{\mathbf{U}}_m)$  and  $\mathbf{U}_1 \in \mathbb{R}^{N_1 \times r}$ ,  $\tilde{\mathbf{U}}_i \in \mathbb{R}^{(N_i-1) \times r}$  for  $i = 2, \dots, m$ .

It can be shown that the normalized problem (4.2) is strongly convex around its global minimizers. For the tensor  $\mathcal{A} \in \mathbb{R}^{N_1 \times N_2 \times \dots \times N_m}$  with  $N_1 \geq N_2 \geq \dots \geq N_m$ , we define the largest rank  $r_m$  such that the problem (4.2) is locally strongly convex. Let

$$(4.3) \quad r_3 := \tilde{N}_{m-2} \lfloor \frac{\tilde{N}_{m-1} \tilde{N}_m}{\tilde{N}_{m-2} + \tilde{N}_{m-1} + \tilde{N}_m - 2} \rfloor,$$

where  $\tilde{N}_{m-2}$ ,  $\tilde{N}_{m-1}$  and  $\tilde{N}_m$  are the largest integers such that (i)  $\tilde{N}_{m-2}$  is even, (ii)  $\tilde{N}_{m-2} \geq \tilde{N}_{m-1} \geq \tilde{N}_m$ , and (iii)  $N_i \geq \tilde{N}_i$  for  $i = m-2, m-1, m$ . Then, the upper bound  $r_m$  is computed recursively by

$$(4.4) \quad r_k := N_{m-k+1} \min\{r_{k-1}, \lfloor \frac{N_{m-k+2} \dots N_m}{N_{m-k+1} + \dots + N_m - k + 1} \rfloor\}, \quad k = 4, \dots, m.$$

The upper bound  $r_m$  is around  $\frac{N_1 \dots N_m}{N_1 + \dots + N_m - m + 1}$  when  $N_1, \dots, N_m$  are large.

The locally strong convexity holds generically when  $r \leq r_m$ . We say a property is generic if it is true on the whole space except a subset with zero measure [14]. The rigorous result is presented in Theorem 4.1.

**Theorem 4.1.** *Suppose that  $N_1 \geq N_2 \geq \dots \geq N_m$  and  $r \leq r_m$  for  $r_m$  in (4.4). Let*

$$\mathcal{A} := \llbracket \mathbf{U}_1^*, [\mathbf{1}^T; \tilde{\mathbf{U}}_2^*], \dots, [\mathbf{1}^T; \tilde{\mathbf{U}}_m^*] \rrbracket,$$

where  $\mathbf{U}_1^* \in \mathbb{R}^{N_1 \times r}$ ,  $\tilde{\mathbf{U}}_i^* \in \mathbb{R}^{(N_i-1) \times r}$ ,  $2 \leq i \leq m$ . Then, for generic  $\tilde{\Theta}^* := (\mathbf{U}_1^*, \dots, \tilde{\mathbf{U}}_m^*)$ , the Hessian matrix  $\nabla^2 \tilde{f}(\tilde{\Theta}^*)$  is positive definite.

*Proof.* See the proof in the Appendix A. ■

Problem (4.2) scales the leading rows to all one vectors, which simplifies the theoretical analysis. In practice, the leading rows can be scaled to arbitrary non-zero vectors. Consider the problem

$$(4.5) \quad \min_{\tilde{\Theta}} \tilde{f}_A(\tilde{\Theta}) := \|\mathcal{A} - \llbracket \mathbf{U}_1, [\mathbf{u}_2^T; \tilde{\mathbf{U}}_2], \dots, [\mathbf{u}_m^T; \tilde{\mathbf{U}}_m] \rrbracket\|_{\mathbb{F}}^2,$$

where  $\mathbf{u}_i \in \mathbb{R}^r$  and  $(\mathbf{u}_i)_j \neq 0, \forall j \in [r]$  and  $i \in [m]$ . Problem (4.2) can be converted Problem (4.5) via some invertible transformations. The invertibility preserves the positive definiteness of the Hessian. Therefore, Problem (4.5) preserves the locally strong convexity.

**Corollary 4.2.** *Suppose  $N_1 \geq N_2 \geq \dots \geq N_m$  and  $r \leq r_m$  in (4.4). Let  $\mathbf{u}_2, \dots, \mathbf{u}_m$  be vectors in  $\mathbb{R}^r$  whose elements are all nonzero and*

$$\mathcal{A} := \llbracket \mathbf{U}_1^*, [\mathbf{u}_2^T; \tilde{\mathbf{U}}_2^*], \dots, [\mathbf{u}_m^T; \tilde{\mathbf{U}}_m^*] \rrbracket,$$

where  $\mathbf{U}_1^* \in \mathbb{R}^{N_1 \times r}$ ,  $\tilde{\mathbf{U}}_i^* \in \mathbb{R}^{(N_i-1) \times r}, 2 \leq i \leq m$ . Then, for generic  $\tilde{\Theta}^* := (\mathbf{U}_1^*, \dots, \tilde{\mathbf{U}}_m^*)$ , the Hessian  $\nabla^2 \tilde{f}_A(\tilde{\Theta}^*)$  is positive definite and there exists an open set  $O$  containing  $\tilde{\Theta}^*$  and a constant  $\lambda > 0$  such that the function  $\tilde{f}_A(\tilde{\Theta})$  in (4.5) is  $\lambda$ -strongly convex in  $O$ .

*Proof.* To simplify the descriptions, here we regard  $\tilde{\Theta}$  as a vector including all optimization variables. There exists a nonsingular matrix  $\mathbf{D}$  such that  $\tilde{f}_A(\tilde{\Theta}) = \tilde{f}(\mathbf{D}\tilde{\Theta})$ . The Hessian  $\nabla^2 \tilde{f}(\mathbf{D}\tilde{\Theta}^*)$  is positive definite by Theorem 4.1. It holds that

$$\nabla^2 \tilde{f}_A(\tilde{\Theta}^*) = \mathbf{D}(\nabla^2 \tilde{f}(\mathbf{D}\tilde{\Theta}^*))\mathbf{D}^T.$$

Thus, the Hessian  $\nabla^2 \tilde{f}_A(\tilde{\Theta}^*)$  is positive definite. Since the eigenvalues of a matrix are continuous with respect to all matrix elements [28], there exists a constant  $\lambda > 0$  and an open set  $O$  containing  $\tilde{\Theta}^*$  such that the smallest eigenvalue of  $\nabla^2 \tilde{f}_A(\tilde{\Theta})$  is not less than  $\lambda$  in  $O$ . In other words,  $\tilde{f}_A(\tilde{\Theta})$  is  $\lambda$ -strongly convex in  $O$ . ■

Based on Corollary 4.2, we can prove that Algorithm 3.1 has a local convergence rate in  $O$  after switching to mixed-precision SGD.

**4.2. Convergence of Algorithm 3.1.** We show that (1) the mixed-precision SignSGD in Algorithm 3.1 converges to a stationary point up to some noise, (2) the mixed-precision SGD in Algorithm 3.1 has a linear convergence rate around the true CP decomposition.

Let  $\Theta^s := (\mathbf{U}_1^s, \dots, \mathbf{U}_m^s)$  denote the factor matrices at the  $s$ -th iteration. Suppose that  $\{\Theta^s\}_{s=0}^{S_1}$  and  $\{\Theta^s\}_{s=S_1+1}^{S_2}$  are generated by mixed-precision SignSGD and mixed-precision SGD respectively in Algorithm 3.1. Recall that  $f$  is the objective function defined in (3.1). We make the following assumptions.

**Assumption 4.3.** *Assume that  $\|\nabla^2 f(\Theta^s)\|_2 \leq L, s = 0, \dots, S_2$ .*

**Assumption 4.4.** *Let  $\tilde{\mathbf{g}}^s$  be the stochastic gradient at the  $s$ -th iteration. Assume that for  $s \in [S_2]$  and  $i \in [m]$ , it holds*

$$\mathbb{E}(\|\tilde{\mathbf{g}}_i^s(j, k) - \mathbf{g}_i^s(j, k)\|^2) \leq \sigma_g^2, \mathbb{E}(\|Q(\tilde{\mathbf{g}}_i^s)(j, k) - \tilde{\mathbf{g}}_i^s(j, k)\|^2) \leq \sigma_Q^2, \forall j \in [N_i], k \in [r].$$

Assumption 4.3 assumes the Hessian matrices are bounded, which is widely used in the convergence analysis of SGD methods. Assumption 4.4 ensures the variance of the stochastic gradient and the quantization error are both bounded. Under Assumption 4.4, the quantized stochastic gradient can be bounded as

$$\begin{aligned} \mathbb{E}(\|\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k) - \mathbf{g}_i^s(j, k)\|) &\leq \mathbb{E}(\|\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k) - \tilde{\mathbf{g}}_i^s(j, k)\|) + \mathbb{E}(\|\tilde{\mathbf{g}}_i^s(j, k) - \mathbf{g}_i^s(j, k)\|) \\ &\leq \sqrt{\mathbb{E}(\|\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k) - \tilde{\mathbf{g}}_i^s(j, k)\|^2)} + \sqrt{\mathbb{E}(\|\tilde{\mathbf{g}}_i^s(j, k) - \mathbf{g}_i^s(j, k)\|^2)} \\ &\leq \sigma_Q + \sigma_g. \end{aligned}$$

**4.2.1. Convergence of Mixed-Precision SignSGD.** We show the convergence of the mixed-precision SignSGD in Algorithm 3.1. Our proof is partially motivated by [4].

**Theorem 4.5.** *Let  $\{\Theta^s\}_{s=0}^{S_1}$  be the sequence generated by the SignSGD update in Algorithm 3.1. Under Assumption 4.3 and Assumption 4.4, we have*

$$(4.6) \quad \sum_{s=0}^{S_1-1} \sum_{i=1}^m \frac{n_i}{N_i} \alpha_s \|\mathbf{g}_i^s\|_1 \leq f(\Theta^0) + \frac{1}{2} r L \sum_{s=0}^{S_1-1} \sum_{i=1}^m n_i \alpha_s^2 + 2r(\sigma_g + \sigma_Q) \sum_{s=1}^{S_1-1} \sum_{i=1}^m n_i \alpha_s.$$

*Proof.* Under Assumption 4.3, it holds that

$$\begin{aligned} (4.7) \quad f(\Theta^{s+1}) &\leq f(\Theta^s) + \sum_{i=1}^m \text{tr}((\mathbf{g}_i^s)^T (\mathbf{U}_i^{s+1} - \mathbf{U}_i^s)) + \sum_{i=1}^m \frac{L}{2} \|\mathbf{U}_i^{s+1} - \mathbf{U}_i^s\|_F^2 \\ &= f(\Theta^s) - \alpha_s \sum_{i=1}^m \text{tr}((\mathbf{g}_i^s(\mathcal{I}_i^s, :))^T \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(\mathcal{I}_i^s, :))) + \sum_{i=1}^m \frac{L}{2} \|\alpha_s \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(\mathcal{I}_i^s, :))\|^2 \\ &= f(\Theta^s) + \frac{1}{2} \alpha_s^2 r L \sum_{i=1}^m n_i - \alpha_s \sum_{i=1}^m \|\mathbf{g}_i^s(\mathcal{I}_i^s, :)\|_1 \\ &\quad + 2\alpha_s \sum_{i=1}^m \sum_{j \in \mathcal{I}_i^s} \sum_{k=1}^r |\mathbf{g}_i^s(j, k)| I(\text{sign}(\mathbf{g}_i^s(j, k)) \neq \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s(j, k)))) , \end{aligned}$$

where  $I$  is the indicator function such that  $I(\text{true}) = 1, I(\text{false}) = 0$ . Considering the part  $I(\text{sign}(\mathbf{g}_i^s(j, k)) \neq \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s(j, k))))$  in the above, we have

$$\begin{aligned} \mathbb{E}[I(\text{sign}(\mathbf{g}_i^s(j, k)) \neq \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k)))] &= \mathbb{P}[I(\text{sign}(\mathbf{g}_i^s(j, k)) \neq \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k)))] \\ &\leq \mathbb{P}[\|\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k) - \mathbf{g}_i^s(j, k)\| \geq |\mathbf{g}_i^s(j, k)|] \\ &\leq \frac{\mathbb{E}[\|\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k) - \mathbf{g}_i^s(j, k)\|]}{|\mathbf{g}_i^s(j, k)|} \\ &\leq \frac{\sigma_g + \sigma_Q}{|\mathbf{g}_i^s(j, k)|} \end{aligned}$$

It implies that

$$\mathbb{E}\left[\sum_{j \in \mathcal{I}_i^s} \sum_{k=1}^r |\mathbf{g}_i^s(j, k)| I(\text{sign}(\mathbf{g}_i^s(j, k)) \neq \text{sign}(\mathbf{Q}(\tilde{\mathbf{g}}_i^s)(j, k)))\right] \leq \sum_{j \in \mathcal{I}_i^s} \sum_{k=1}^r (\sigma_g + \sigma_Q) = n_i r (\sigma_g + \sigma_Q).$$

Then, we take expectation on both sides of (4.7) and get

$$f(\Theta^{s+1}) \leq f(\Theta^s) + \frac{1}{2}\alpha_s^2 rL \sum_{i=1}^m n_i - \alpha_s \sum_{i=1}^m \frac{n_i}{N_i} \|\mathbf{g}_i^s\|_1 + 2\alpha_s r(\sigma_g + \sigma_Q) \sum_{i=1}^m n_i.$$

After summing up both sides for  $s = 0, \dots, S_1 - 1$  and rearrangement, we have

$$\sum_{s=0}^{S_1-1} \sum_{i=1}^m \frac{n_i}{N_i} \alpha_s \|\mathbf{g}_i^s\|_1 \leq f(\Theta^0) - f(\Theta^{S_1}) + \frac{1}{2}rL \sum_{s=0}^{S_1-1} \sum_{i=1}^m n_i \alpha_s^2 + 2r(\sigma_g + \sigma_Q) \sum_{s=0}^{S_1-1} \sum_{i=1}^m n_i \alpha_s.$$

It implies the result (4.6) since  $f(\Theta^{S_1}) \geq 0$ .  $\blacksquare$

In practice, we usually choose a relatively large constant learning rate  $\alpha_s = \alpha$  to accelerate the convergence at the beginning. We prove in Corollary 4.6 that a constant learning rate provides  $O(\frac{1}{T_1})$  convergence rate up to some noise.

**Corollary 4.6.** *Under conditions of Theorem 4.5, if  $\alpha_s = \alpha$ , then*

$$\min_{s=0, \dots, S_1-1} \|\mathbf{g}_i^s\|_1 \leq O\left(\frac{1}{S_1}\right) + \frac{1}{\gamma} \left(\frac{\alpha L}{2} + 2\sigma_g + 2\sigma_Q\right) r \sum_{i=1}^m n_i,$$

where  $\gamma = \min_{i \in [m]} \frac{n_i}{N_i}$ .

*Proof.* Equation (4.6) implies that

$$\begin{aligned} \min_{s=0, \dots, S_1-1} \|\mathbf{g}_i^s\|_1 &\leq \frac{1}{\gamma} \frac{f(\Theta^0)}{\sum_{s=0}^{S_1-1} \alpha_s} + \frac{1}{\gamma} \left( \frac{1}{2} rL \sum_{i=1}^m n_i \frac{\sum_{t=0}^{T_1-1} \alpha_s^2}{\sum_{t=0}^{T_1-1} \alpha_s} + 2r(\sigma_g + \sigma_Q) \sum_{i=1}^m n_i \right) \\ &= \frac{f(\Theta^0)}{S_1 \gamma \alpha} + \frac{1}{\gamma} \left( \frac{\alpha L}{2} + 2\sigma_g + 2\sigma_Q \right) r \sum_{i=1}^m n_i \\ &= O\left(\frac{1}{S_1}\right) + \frac{1}{\gamma} \left( \frac{\alpha L}{2} + 2\sigma_g + 2\sigma_Q \right) r \sum_{i=1}^m n_i. \end{aligned}$$

**4.2.2. Convergence of Mixed-Precision SGD.** In this subsection, we show the locally linear convergence rate of the mixed-precision SGD in Algorithm 3.1. We make the following extra assumption.

**Assumption 4.7.** *Assume that for some  $\theta \in [0, 1)$ , it holds*

$$\|\mathbb{E}(Q(\tilde{\mathbf{g}}_i^s) - \tilde{\mathbf{g}}_i^s)\| \leq \theta \|\mathbf{g}_i^s\|, \quad i \in [m], s = S_1 + 1, \dots, S_2.$$

Assumption 4.7 assumes the quantized stochastic gradient is a good descent direction in expectation. If the quantization function  $Q$  uses independent stochastic rounding, then Assumption 4.7 is true for  $\theta = 0$  since  $\mathbb{E}(Q(\tilde{\mathbf{g}}_i^s)) = \mathbb{E}(\tilde{\mathbf{g}}_i^s) = \mathbf{g}_i^s$ . Assumption 4.7 still holds for deterministic rounding as long as the quantization error is not large.

Let  $\mathbf{u}_i := \mathbf{U}_i^{S_1}(1, \cdot) \in \mathbb{R}^r$  for  $i = 2, \dots, m$ . Suppose  $\mathcal{A} := [\mathbf{U}_1^*, [\mathbf{u}_2^T; \tilde{\mathbf{U}}_2^*], \dots \circ [\mathbf{u}_m^T; \tilde{\mathbf{U}}_m^*]]$ , where  $\mathbf{U}_1^* \in \mathbb{R}^{N_1 \times r}$ ,  $\tilde{\mathbf{U}}_i^* \in \mathbb{R}^{(N_i-1) \times r}$ ,  $2 \leq i \leq m$ . The objective function now becomes  $\tilde{f}_{\mathcal{A}}(\tilde{\Theta})$  as in (4.5), which is locally  $\lambda$ -strongly convex by Corollary 4.2. Consequently, the convergence result of SGD for strongly convex functions can be applied.

**Theorem 4.8.** Suppose the tensor  $\mathcal{A}$  satisfies the conditions of Corollary 4.2. Let  $\{\Theta^s\}_{s=S_1}^{S_2}$  be the sequence generated by the mixed-precision SGD in Algorithm 3.1 and  $\alpha_s = \alpha$  be the learning rate. If  $\tilde{\Theta}^s$  is in the set  $O$  as in Corollary 4.2 and  $\mathbb{E}(\|Q(\tilde{\mathbf{g}}^s)\|) \leq G$  for  $s = S_1, \dots, S_2$ , then under Assumption 4.3, Assumption 4.4, and Assumption 4.7, it holds that

$$\mathbb{E}(\|f(\Theta^{S_2})\|^2) \leq \frac{\alpha LG}{2\lambda(1-\theta)} + (1 - \alpha\lambda(1-\theta))^{S_2-S_1} (f(\Theta^{S_1}) - \frac{\alpha LG}{2\lambda(1-\theta)})$$

where  $\lambda$  is the strong convexity parameter in Corollary 4.2.

*Proof.* Under Assumption 4.7, the gradient  $\tilde{\mathbf{g}}_i^s$  satisfies

$$\begin{aligned} \mathbb{E}((\mathbf{g}_i^s)^T Q(\tilde{\mathbf{g}}_i^s)) &= (\mathbf{g}_i^s)^T \mathbb{E}(\tilde{\mathbf{g}}_i^s) + (\mathbf{g}_i^s)^T \mathbb{E}(Q(\tilde{\mathbf{g}}_i^s) - \tilde{\mathbf{g}}_i^s) \\ &\geq \|\mathbf{g}_i^s\|^2 - \|\mathbf{g}_i^s\| \mathbb{E}\|Q(\tilde{\mathbf{g}}_i^s) - \tilde{\mathbf{g}}_i^s\| \\ &\geq (1-\theta)\|\mathbf{g}_i^s\|^2. \end{aligned}$$

By Corollary 4.2, the function  $\tilde{f}_A(\tilde{\Theta})$  is  $\lambda$ -strongly convex in  $O$  containing  $\tilde{\Theta}^*$ . Algorithm 3.1 is minimizing the function  $f_A(\Theta)$  after switching to mixed-precision SGD. It also holds that  $f(\Theta^{S_2}) = f_A(\tilde{\Theta}^{S_2})$ . Therefore, the result is a direct conclusion of Theorem 4.6 in [7].  $\blacksquare$

## 5. Numerical Experiments.


**5.1. Implementation Details.** Recall that the block stochastic gradient is computed as in Algorithm 3.2. The quantization function  $Q_1$  use  $\text{FP}_{16}$ , scale factor  $\delta = 1$ , and deterministic rounding, i.e.,  $Q_1 = Q_{\text{FP}_{16},1}^D$ . The quantization function  $Q_2$  use  $\text{INT}_b$  precision and deterministic rounding. When quantizing the matrix  $\mathbf{X}$ , we use the scale factor  $\delta = \frac{\max \|\mathbf{X}\|}{c}$ , where  $c \geq 2^{b-1} - 1$ . Specifically, we use  $c = 10, 30, 200$  for  $\text{INT}_2, \text{INT}_4, \text{and } \text{INT}_8$  respectively. In this section, the precision of Algorithm 3.1 always means the precision of  $Q_2$ .

Our implementation uses the Python package CuPy [43]. For fair comparisons between different precisions, we implement the matrix multiplication by CUTLASS kernels [32]. However, due to the lack of support for extremely low-bit fixed-point integer representations in Python, we only compare the running time between  $\text{INT}_8$  and  $\text{FP}_{32}$  on GPU. The learning rate for the mixed-precision SignSGD in Algorithm 3.1 is set as 0.5 initially and is updated as  $\alpha = 0.3\alpha$  every 1000 iterations. The mixed-precision SGD stage uses the constant learning rate  $\alpha = 0.01$ . For tensors of orders 3, 4, 5, we use the sample sizes  $|\mathcal{I}_i| = 0.2N_i, 0.3N_i, 0.4N_i$  respectively. The size of the sampled sub-tensor is roughly 1% of the original tensor.

Suppose that Algorithm 3.1 outputs the factor matrices  $\{\mathbf{U}_i\}_{i=1}^m$  for the input tensor  $\mathcal{A}$ . We use a relative error to measure the quality of our results, which is defined as

$$\text{error} = \frac{\|\mathcal{A} - \llbracket \mathbf{U}_1, \dots, \mathbf{U}_m \rrbracket\|_F}{\|\mathcal{A}\|_F}.$$

**5.2. Synthetic Examples.** We first test the runtime and convergence of Algorithm 3.1 under various precisions on some synthetic tensor benchmarks.



**Figure 1.** Performance of Algorithm 3.1 with and without SignSGD initialization.

Dimension	Sample size	Rank	INT <sub>8</sub> time (s)	FP <sub>32</sub> time (s)	Speed up
(1200,1200,1200)	(240,240,240)	256	17.96	28.15	<b>1.56</b> ×
(200,200,200,200)	(60,60,60,60)	64	29.90	56.44	<b>1.88</b> ×
(60,60,60,60,60)	(24,24,24,24,24)	32	41.71	114.89	<b>2.75</b> ×

**Table 4**

Time comparison between FP<sub>32</sub> and INT<sub>8</sub> of Algorithm 3.1 for various dimensions

**5.2.1. Role of SignSGD Initialization.** This section runs the experiment in full precision to show the influence of SignSGD initialization to the convergence of the whole algorithm. The results with different initialization methods are shown in Figure 1. The “max” in Figure 1 is the maximum absolute value of each  $\{\mathbf{U}_i\}_{i=1}^m$ . Algorithm 3.1 without SignSGD is trapped by a stationary point and fails to converge with max = 1.0 as shown in Figure 1a. After we decrease max to 0.1 and 0.01, Algorithm 3.1 without SignSGD stays at zero, which is a stationary point. In contrast, Algorithm 3.1 with SignSGD converges well for max=1.0,0.1,0.01. The result demonstrates that the SignSGD initialization can greatly improve the convergence of Algorithm 3.1.

**5.2.2. Time Comparison in Different Precisions.** We test the runtime of Algorithm 3.1 to reach the same relative error  $10^{-3}$  under different precisions. We specifically compare the runtime of Algorithm 3.1 with INT<sub>8</sub> and FP<sub>32</sub> respectively, and the result is summarized in Table 4. Figure 2a, 2b, 2c show that the runtime increases linearly as the tensor rank increases for both low-precision and full-precision. The reduction ratio remains the same as the rank changes. We can observe significant time savings when using the INT<sub>8</sub> format for all sizes, ranks, and orders. The time saving is also more remarkable as the tensor order increases. This is because  $m$  large-size matrix multiplications are computed in low precision in Algorithm 3.1 for tensors with order  $m$ . Therefore, higher order  $m$  brings in more time savings. The detailed complexity analysis is in Section 3.2.





Figure 2. Time comparison between  $FP_{32}$  and  $INT_8$  of Algorithm 3.1 for various ranks

**5.2.3. Convergence Comparison in Different Precisions.** We further evaluate the convergence of Algorithm 3.1 under various precisions. We compare precisions  $INT_2$ ,  $INT_4$ ,  $INT_8$ , and  $FP_{32}$ , where the computation of  $INT_2$  and  $INT_4$  is simulated by  $FP_{32}$ . The simulation simply rounds the scaled number into the nearest integer and then clamps it into the representation range. The convergence of  $INT_4$  and  $INT_8$  precision are almost the same as the convergence of  $FP_{32}$ , so they are combined in Figure 3. The final relative error of low-precision Algorithm 3.1 is slightly worse than the full-precision version due to the quantization error. The quantization error also causes the slow convergence for rank 200 and the divergence for higher ranks of  $INT_2$  precision. The noisy ball term in Corollary 4.6 for SignSGD depends on the rank  $r$  and the quantization error  $\sigma_Q$ . Therefore, a large rank  $r$  and large quantization error  $\sigma_Q$  may lead to bad convergence due to the large noisy ball. The mixed-precision SGD part starts at around the 2000th iteration. Figure 3 shows that the mixed-precision SGD has a linear convergence rate which matches the theoretical result in Theorem 4.8. The slower convergence of the mixed-precision SGD part of  $INT_2$  precision in Figure 3a is caused by the large  $\theta$  in Assumption 4.7 due to the quantization error.

### 5.3. Real Datasets.

**5.3.1. Coil-100 Dataset.** The Coil-100 dataset [38] contains the images of 100 objects in 72 different poses. Each image has size  $128 \times 128 \times 3$ , where  $128 \times 128$  is the number of pixels and 3 represents the 3 RGB channels. Thus, the size of the formed tensor  $\mathcal{A}$  is  $128 \times 128 \times 3 \times 7200$ . The CP decomposition is applied for dimension reduction. The fourth-factor matrix  $\mathbf{U}_4 \in \mathbb{R}^{7200 \times r}$  can be used as features for clustering and classification tasks.

We run Algorithm 3.1 on the tensor  $\mathcal{A}$  with precisions  $INT_2$ ,  $INT_4$ ,  $INT_8$ , and  $FP_{32}$ . The test employs rank  $r = 16$  and sample size  $(32, 32, 3, 1440)$ . As shown in Figure 4, the convergence trends are similar for all precisions. Higher numerical precisions produce smaller relative errors in the final solution, but the difference is insignificant. All of our relative errors are better than the best reported result in [3], which is 0.314. Regarding the running time,  $INT_8$  takes 11 seconds for 1000 iterations while  $FP_{32}$  takes 31 seconds. Algorithm 3.1 of the  $INT_8$  precision is **2.8** times faster. The experiment demonstrates that our proposed mixed-precision algorithm can effectively reduce the computation cost of CP tensor decomposition on real-world datasets with negligible accuracy loss.



**Figure 3.** Convergence comparison of various precision on the tensor of dimension  $(1200,1200,1200)$ .


**5.3.2. MRI Dataset.** Magnetic Resonance Imaging (MRI) is widely used in brain science and clinic diagnosis. Low-rank tensor decomposition can be applied to denoise practical MR images [52]. This experiment uses the data from the NYU fastMRI Initiative database [34, 53]. The original data is in a Fourier space and forms a complex tensor  $\mathcal{K}$  of size  $16 \times 640 \times 320$ . The tensor  $\mathcal{K}$  in real-world is typically corrupted by noises. In our test, we intentionally corrupt  $\mathcal{K}$  by the noise tensor  $\mathcal{N}$ . The real part and imaginary part of  $\mathcal{N}$  both obey the normal distribution with mean 0 and variance  $\tau^2$ . Let the corrupted tensor be  $\hat{\mathcal{K}} = \mathcal{K} + \mathcal{N}$ , then the inverse Fourier transform is applied to  $\hat{\mathcal{K}}$  to get  $\hat{\mathcal{A}} \in \mathbb{C}^{16 \times 640 \times 320}$ . Next, we use Algorithm 3.1 to find a low-rank approximation  $\mathcal{A}$  of the noisy tensor  $\hat{\mathcal{A}}$  for noise removal. Finally, the gray-scale image  $\mathbf{M} \in \mathbb{R}^{640 \times 320}$  is reconstructed as

$$\mathbf{M}_{i,j} = \sqrt{\sum_{k=1}^{16} |\mathcal{A}(k, i, j)|^2}.$$

The image  $\mathbf{M}$  is further cropped into size  $320 \times 320$  by selecting  $\mathbf{M}(161 : 480, :)$ . Let  $\mathbf{M}_{\text{truth}}$  be the ground-truth of the image, then the relative error is computed as

$$\text{error} = \frac{\|\mathbf{M} - \mathbf{M}_{\text{truth}}\|}{\|\mathbf{M}_{\text{truth}}\|}.$$





**Figure 4.** Converge curves of Algorithm 3.1 in various precisions on Coil-100 data

$\tau$	w/o CP	INT <sub>4</sub>	INT <sub>8</sub>	FP <sub>32</sub>
1.0e-5	0.379	0.115	0.109	0.109
1.5e-5	0.639	0.173	0.164	0.165
2.0e-5	0.913	0.246	0.236	0.236


**Table 5**

Relative errors before and after removing noises of MRI by Algorithm 3.1

This experiment applied Algorithm 3.1 to the complex tensors. We would like to remark that Algorithm 3.1 is designed for real tensors, but we can extend the algorithm to complex tensors by considering the real part and the imaginary part separately.

Our experiment uses rank  $r = 200$ , sample size  $(16, 64, 64)$  and noise tensors  $\mathcal{N}$  of standard deviations  $\tau = 1.0 \times 10^{-5}, 1.5 \times 10^{-5}, 2.0 \times 10^{-5}$ , respectively. We test the performance of Algorithm 3.1 in various precisions. The test results are presented in Table 5 and Figure 5. Table 5 lists the relative errors before and after CP decompositions. Algorithm 3.1 successfully removes the noises and reduces the errors. Moreover, the performance of Algorithm 3.1 in INT<sub>4</sub> and INT<sub>8</sub> is similar to FP<sub>32</sub>. Figure 5 compares the recovered images and the ground-truth image. The images obtained via INT<sub>4</sub>, INT<sub>8</sub>, FP<sub>32</sub> mixed-precision CP decomposition are visually identical and vastly superior to those without noise removal. The results demonstrate that the mixed-precision Algorithm 3.1 is capable of producing accurate decomposition for noisy MRI datasets.

**5.4. FPGA Demonstration for Edge Computing.** The proposed mixed-precision CP decomposition can reduce the computing cost on both cloud and edge devices. Here we implement Algorithm 3.1 on a Field Programmable Gate Array (FPGA) to demonstrate its benefit on resource-constrained edge devices. FPGAs are widely used for edge computing due to their energy efficiency, flexible reconfigurability, and fast time-to-market [48]. However, FP-



**Figure 5.** Error images before and after removing noises of MRI by Algorithm 3.1

Precision	Time (s)	BRAM	FF	Power (W)	Energy (J)
INT <sub>2</sub>	<b>3.08</b>	<b>12</b>	<b>86025</b>	<b>8.4</b>	<b>25.87</b>
INT <sub>8</sub>	<b>3.08</b>	15	197868	9.5	29.26
FP <sub>32</sub>	8.61	250	624416	13.3	114.4

**Table 6**

FPGA implementations of Algorithm 3.1 in different precisions

GAs have very limited memory and computing resources, therefore, it is desired to use low numerical precision to save the hardware cost in massive engineering applications.

We consider a rank-20 tensor with size (100, 100, 100) and sample size (20, 20, 20). Table 6 shows the performance of Algorithm 3.1 on FPGA in different precisions. The integer operations accelerate Algorithm 3.1 about **2.8** times compared to FP<sub>32</sub>. The running time of Algorithm 3.1 in INT<sub>2</sub> and INT<sub>8</sub> are dominated by higher precision parts in the algorithm. Consequently, INT<sub>2</sub> and INT<sub>8</sub> have the similar running time. The usage of BRAM (block random-access memory) of INT<sub>2</sub> and INT<sub>8</sub> is about **20** times less than FP<sub>32</sub>. Due to the reduced hardware resource requirements, the power consumption of INT<sub>2</sub> and INT<sub>8</sub> is also reduced. The energy cost of FP<sub>32</sub> is 4 times more than the energy consumed by INT<sub>2</sub> and INT<sub>8</sub>. The number of FF (Flip-Flop) used by INT<sub>2</sub> is **2.3** times less than INT<sub>8</sub> and **7.2** times less than FP<sub>32</sub>, respectively. In summary, the reduction of time and resources on FPGAs successfully demonstrates the effectiveness of our proposed mixed-precision algorithm on resource-constrained devices.

**6. Conclusion.** This paper has proposed a mixed-precision stochastic gradient method for the CP tensor decomposition problem. First, the stochastic gradient is computed in mixed-precision to reduce the runtime and computation cost. Then, we develop a two-stage optimization algorithm to solve the CP decomposition problem using mixed-precision gradients. The convergence of the proposed algorithm has been proved. We have shown that the CP decomposition problem is locally strongly convex after proper normalization. Consequently, the mixed-precision SGD stage in our algorithm can have a linear convergence rate. A set of numerical experiments on GPUs and on edge devices have successfully demonstrated that our mixed-precision algorithm can significantly reduce the computation costs and latency compared to the full-precision algorithm while maintaining high accuracy. The proposed mixed-precision stochastic gradient method can be applied to many gradient-based CP decomposition algorithms. It will be an interesting future topic to study the applications of mixed-precision gradients to other optimization algorithms.

**Appendix A. Proof of Theorem 4.1.** In the appendix, we will give the concrete proof of Theorem 4.1. We first show the Hessian is positive definite if and only if the Jacobian matrix has full column rank. Then, we prove that the Jacobian matrix has a full column rank in general.

We consider the following vector-valued function

$$(A.1) \quad \tilde{\mathbf{h}}(\tilde{\Theta}) = \text{vec} \left( [\mathbf{U}_1, [\mathbf{1}^T; \tilde{\mathbf{U}}_2], \dots, [\mathbf{1}^T; \tilde{\mathbf{U}}_m]] \right),$$

where  $\text{vec}(\cdot)$  is the vectorization function. It holds that  $\tilde{f}(\tilde{\Theta}) = \|\tilde{\mathbf{h}}(\tilde{\Theta}) - \text{vec}(\mathcal{A})\|^2$ .

**Lemma A.1.** *Suppose that  $\tilde{f}(\tilde{\Theta}) = 0$ , then  $\nabla^2 \tilde{f}(\tilde{\Theta})$  is positive definite if and only if the Jacobian  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank.*

*Proof.*  $\tilde{f}(\tilde{\Theta}) = 0$  implies that  $\tilde{\mathbf{h}}(\tilde{\Theta}) - \text{vec}(\mathcal{A}) = \mathbf{0}$ . Then, we have

$$\nabla \tilde{f}(\tilde{\Theta}) = 2(\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta}))^T \tilde{\mathbf{h}}(\tilde{\Theta}), \quad \nabla^2 \tilde{f}(\tilde{\Theta}) = 2(\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta}))^T \mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta}).$$

As a result,  $\nabla^2 \tilde{f}(\tilde{\Theta})$  is positive definite if and only if  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank. ■

In the following, we first prove  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  generally has full column rank for the order  $m = 3$  and the even  $N_1$ , then we extend the result to general orders and dimensions.

**Proposition A.2.** *When  $m = 3$ ,  $N_1$  is even,  $N_1 \geq N_2 \geq N_3 \geq 3$ , and  $r \leq N_1 \lfloor \frac{n_2 n_3}{N_1 + N_2 + N_3 - 2} \rfloor$ , the Jacobian matrix  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank for generic  $\tilde{\Theta} = (\mathbf{U}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3)$ .*

*Proof.* It suffices to only consider  $r = N_1 \lfloor \frac{n_2 n_3}{N_1 + N_2 + N_3 - 2} \rfloor$ . The Jacobian matrix  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank if and only if  $(\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta}))^T \mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has non-zero determinant, which is a polynomial function  $p$  in terms of variables  $\tilde{\Theta}$ . The conclusion is then equivalent to that  $p(\tilde{\Theta})$  is nonzero for generic  $\tilde{\Theta}$ . Thus, it suffices to show  $p(\tilde{\Theta})$  is not the constant 0 polynomial [14]. In the following, we will construct the specific  $\tilde{\Theta}$  such that  $p(\tilde{\Theta}) \neq 0$ .

Here we denote  $\mathbf{U}_i = [\mathbf{1}^T; \tilde{\mathbf{U}}_i]$  for convenience. Let  $k = \frac{N_1}{2}$  and  $R = 2 \lfloor \frac{n_2 n_3}{N_1 + N_2 + N_3 - 2} \rfloor$ . We evenly split  $\{1, \dots, r\}$  into  $k$  groups such that each group has  $R = 2 \lfloor \frac{n_2 n_3}{N_1 + N_2 + N_3 - 2} \rfloor$  elements.

We denote the groups by  $\text{group}_1, \dots, \text{group}_k$ . For  $i = 1, \dots, k$ , let

$$\mathbf{J}_i = \left[ \left[ [\mathbf{e}_{2i-1}, \mathbf{e}_{2i}] \otimes \mathbf{U}_2(:, j) \otimes \mathbf{U}_3(:, j) \right]_{j=1}^r, \left[ \mathbf{U}_1(:, j) \otimes \tilde{\mathbf{I}}_{N_2} \otimes \mathbf{U}_3(:, j) \right]_{j \in \text{group}_i}, \right. \\ \left. \left[ \mathbf{U}_1(:, j) \otimes \mathbf{U}_2(:, j) \otimes \tilde{\mathbf{I}}_{N_3} \right]_{j \in \text{group}_i} \right],$$

where  $\tilde{\mathbf{I}}_n := \mathbf{I}_n(2 : n, :)$ . The Jacobian  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  can be written as  $[\mathbf{J}_1, \dots, \mathbf{J}_k]$ .

We construct the matrix  $\mathbf{U}_1$ , where for  $j \in \text{group}_i$

$$\mathbf{U}_1(l, j) = 0, \text{ if } l \neq 2i - 1, 2i.$$

For such  $\mathbf{U}_1$ , the Jacobian  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank if and only if each  $\mathbf{J}_i$  has full column rank. In the following, we will construct  $\tilde{\Theta}$  such that  $\mathbf{J}_1$  has full column rank.

Since  $N_1 \geq N_2 \geq N_3 \geq 3$ , it holds that

$$R = 2 \lfloor \frac{n_2 n_3}{N_1 + N_2 + N_3 - 2} \rfloor \leq \frac{2n_2 n_3}{2n_2 + N_3 - 2} \leq \frac{2n_2 n_3}{2n_2 - 1} < \frac{2n_2 n_3}{2n_2} = N_3.$$

Thus, we have  $R < N_3 \leq N_2$ . Let  $\mathbf{a}_1, \dots, \mathbf{a}_r$  be pairwise independent vectors,  $\mathbf{b}_1, \dots, \mathbf{b}_{N_2}$  and  $\mathbf{c}_1, \dots, \mathbf{c}_{N_3}$  be orthonormal basis of  $\mathbb{R}^{N_2}$  and  $\mathbb{R}^{N_3}$ , respectively. The orthonormal basis can be chosen such that all leading entries are nonzero. Denote the matrix

$$\mathbf{P} := \left[ \left[ \mathbf{I}_2 \otimes \mathbf{b}_j \otimes \mathbf{c}_j, \mathbf{a}_j \otimes \tilde{\mathbf{I}}_{N_2} \otimes \mathbf{c}_j, \mathbf{a}_j \otimes \mathbf{b}_j \otimes \tilde{\mathbf{I}}_{N_3} \right]_{j=1}^R, \mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3 \right],$$

where

$$\mathbf{Q}_1 = [\mathbf{I}_2 \otimes \mathbf{b}_i \otimes \mathbf{c}_j]_{R+1 \leq i \leq N_2, R+1 \leq j \leq N_3},$$

$$\mathbf{Q}_2 = [\mathbf{I}_2 \otimes \mathbf{b}_i \otimes (\mathbf{c}_{2j-1} + \mathbf{c}_{2j})]_{R+1 \leq i \leq N_2, 1 \leq j \leq R/2},$$

$$\mathbf{Q}_3 = [\mathbf{I}_2 \otimes (\mathbf{b}_{2i-1} + \mathbf{b}_{2i}) \otimes \mathbf{c}_j]_{1 \leq i \leq R/2, R+1 \leq j \leq N_3}.$$

Next, we show the matrix  $\mathbf{P}$  has full column rank. It is equivalent to proving that  $\mathbf{P}\mathbf{x} = \mathbf{0} \Leftrightarrow \mathbf{x} = \mathbf{0}$ . We first prove the coefficients for  $\mathbf{Q}_1$  is zero. For some  $R+1 \leq i \leq N_2, R+1 \leq j \leq N_3$ , it holds that

$$(\mathbf{I}_2 \otimes \mathbf{b}_i^T \otimes \mathbf{c}_j^T) \mathbf{P}\mathbf{x} = \mathbf{I}_2 \boldsymbol{\lambda} = \mathbf{0} \Rightarrow \boldsymbol{\lambda} = \mathbf{0},$$

where  $\boldsymbol{\lambda}$  is the coefficient vector corresponding to  $\mathbf{I}_2 \otimes \mathbf{b}_i \otimes \mathbf{c}_j$  in  $\mathbf{x}$ . Thus the coefficient for  $\mathbf{Q}_1$  is zero. Then we show the coefficient for  $\mathbf{Q}_2$  is zero. For some  $R+1 \leq i \leq N_2, 1 \leq j \leq R/2$ , it holds

$$(\mathbf{I}_2 \otimes \mathbf{b}_i^T \otimes \mathbf{c}_{2j-1}^T) \mathbf{P}\mathbf{x} = (\mathbf{a}_{2j-1} \otimes \mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2}, \mathbf{I}_2) \begin{bmatrix} \boldsymbol{\lambda}_1 \\ \boldsymbol{\mu} \end{bmatrix} = \mathbf{0},$$

$$(\mathbf{I}_2 \otimes \mathbf{b}_i^T \otimes \mathbf{c}_{2j}^T) \mathbf{P}\mathbf{x} = (\mathbf{a}_{2j} \otimes \mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2}, \mathbf{I}_2) \begin{bmatrix} \boldsymbol{\lambda}_2 \\ \boldsymbol{\mu} \end{bmatrix} = \mathbf{0},$$

where  $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \boldsymbol{\mu}$  are coefficients corresponding to  $\mathbf{a}_{2j-1} \otimes \tilde{\mathbf{I}}_{N_2} \otimes \mathbf{c}_{2j}$ ,  $\mathbf{a}_{2j} \otimes \tilde{\mathbf{I}}_{N_2} \otimes \mathbf{c}_{2j}$  and  $\mathbf{I}_2 \otimes \mathbf{b}_i \otimes (\mathbf{c}_{2j-1} + \mathbf{c}_{2j})$  respectively. By above equations, we know

$$\boldsymbol{\mu} = -(\mathbf{a}_{2j-1} \otimes \mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2}) \boldsymbol{\lambda}_1 = -(\mathbf{a}_{2j} \otimes \mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2}) \boldsymbol{\lambda}_2.$$

Thus,  $\mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2} \boldsymbol{\lambda}_1 = \mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2} \boldsymbol{\lambda}_2 = \mathbf{0}$  since  $\mathbf{a}_{2j-1}, \mathbf{a}_{2j}$  are linearly independent. It implies that  $\boldsymbol{\mu} = \mathbf{0}$ . Therefore, the coefficient for  $\mathbf{Q}_2$  is zero. Similarly, we can prove the coefficient for  $\mathbf{Q}_3$  is zero by using exactly the same technique.

Next, we show the coefficient of  $\mathbf{a}_j \otimes \tilde{\mathbf{I}}_{N_2} \otimes \mathbf{c}_j$  is zero. For some  $1 \leq i \leq N_2$  and  $i \neq j$ , it holds

$$(\mathbf{I}_2 \otimes \mathbf{b}_i^T \otimes \mathbf{c}_j^T) \mathbf{P} \mathbf{x} = (\mathbf{a}_j \otimes \mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2}, \mathbf{a}_i \otimes \mathbf{c}_j^T \tilde{\mathbf{I}}_{N_3}) \begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\lambda}_j \end{bmatrix} = \mathbf{0} \Rightarrow \mathbf{b}_i^T \tilde{\mathbf{I}}_{N_2} \boldsymbol{\mu} = \mathbf{c}_j^T \tilde{\mathbf{I}}_{N_3} \boldsymbol{\lambda}_j = 0,$$

where  $\boldsymbol{\mu}, \boldsymbol{\lambda}_j$  are coefficients corresponding to  $\mathbf{a}_j \otimes \tilde{\mathbf{I}}_{N_2} \otimes \mathbf{c}_j, \mathbf{a}_i \otimes \mathbf{b}_i \otimes \tilde{\mathbf{I}}_{N_3}$  respectively. The above equation holds for every  $i$  such that  $1 \leq i \leq N_2$  and  $i \neq j$ . Therefore, we have

$$[\tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_1, \dots, \tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_{j-1}, \tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_{j+1}, \dots, \tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_{N_2}]^T \boldsymbol{\mu} = \mathbf{0} \Rightarrow \boldsymbol{\mu} = \mathbf{0}.$$

The above holds because  $[\tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_1, \dots, \tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_{j-1}, \tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_{j+1}, \dots, \tilde{\mathbf{I}}_{N_2}^T \mathbf{b}_{N_2}]^T \in \mathbb{R}^{(N_2-1) \times (N_2-1)}$  is nonsingular. It proves that the coefficient for  $\mathbf{a}_j \otimes \tilde{\mathbf{I}}_{N_2} \otimes \mathbf{c}_j$  is zero. Similarly, we can prove the coefficient for  $\mathbf{a}_j \otimes \mathbf{b}_j \otimes \tilde{\mathbf{I}}_{N_3}$  is zero.

The only remaining part in  $\mathbf{P}$  is  $[\mathbf{I}_2 \otimes \mathbf{b}_j \otimes \mathbf{c}_j]_{j=1}^R$ .  $[\mathbf{I}_2 \otimes \mathbf{b}_j \otimes \mathbf{c}_j]_{j=1}^R$  has full column rank since  $\mathbf{b}_1, \dots, \mathbf{b}_R$  are linearly independent. Thus, the coefficients corresponding to  $[\mathbf{I}_2 \otimes \mathbf{b}_j \otimes \mathbf{c}_j]_{j=1}^R$  are also zero. It finishes the proof that  $\mathbf{P} \mathbf{x} = \mathbf{0} \Leftrightarrow \mathbf{x} = \mathbf{0}$ . Thus,  $\mathbf{P}$  has full column rank.

Let

$$\mathbf{Q} := \begin{bmatrix} \left[ \frac{\mathbf{b}_i}{(\mathbf{b}_i)_1} \otimes \frac{\mathbf{c}_j}{(\mathbf{c}_j)_1} \right]_{R+1 \leq i \leq N_2, R+1 \leq j \leq N_3}, \left[ \frac{\mathbf{b}_i}{(\mathbf{b}_i)_1} \otimes \frac{\mathbf{c}_{2j-1} + \mathbf{c}_{2j}}{(\mathbf{c}_{2j-1})_1 + (\mathbf{c}_{2j})_1} \right]_{R+1 \leq i \leq N_2, 1 \leq j \leq R/2}, \\ \left[ \frac{\mathbf{b}_{2i-1} + \mathbf{b}_{2i}}{(\mathbf{b}_{2i-1})_1 + (\mathbf{b}_{2i})_1} \otimes \frac{\mathbf{c}_j}{(\mathbf{c}_j)_1} \right]_{1 \leq i \leq R/2, R+1 \leq j \leq N_3}. \end{bmatrix}$$

The number of columns of  $\mathbf{Q}$  is

$$c = (N_2 - R)(N_3 - R) + (N_2 - R) \frac{R}{2} + (N_3 - R) \frac{R}{2} = n_2 n_3 - \frac{R}{2} (N_2 + N_3).$$

It holds that,

$$c - (r - R) = n_2 n_3 - \frac{R}{2} (N_2 + N_3 - 2) - \frac{n_1 R}{2} = n_2 n_3 - \frac{R}{2} (N_1 + N_2 + N_3 - 2) \geq 0.$$

Let  $\mathbf{U}_1(1 : 2, j) = \mathbf{a}_j$ ,  $\tilde{\mathbf{U}}_2(:, j) = \mathbf{b}_j(2 :)/(\mathbf{b}_j)_1$ ,  $\tilde{\mathbf{U}}_3(:, j) = \mathbf{c}_j(2 :)/(\mathbf{c}_j)_1$  for  $j = 1, \dots, R$  and  $\tilde{\mathbf{U}}_2(:, j), \tilde{\mathbf{U}}_3(:, j)$  be vectors such that

$$\left[ [1; \tilde{\mathbf{U}}_2(:, j)] \otimes [1; \tilde{\mathbf{U}}_3(:, j)] \right]_{R+1 \leq j \leq r} = \mathbf{Q}(:, 1 : r - R).$$

Then, all columns of  $\mathbf{J}_1$  are from  $\mathbf{P}$ . We have shown that  $\mathbf{P}$  has full column rank, so  $\mathbf{J}_1$  must have full column rank. The same technique can be applied to  $\mathbf{J}_1, \dots, \mathbf{J}_k$ .

Now we have proven that  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\boldsymbol{\Theta}})$  has full column rank for some  $\tilde{\boldsymbol{\Theta}}$ . Therefore, the Jacobian  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\boldsymbol{\Theta}})$  has full column rank for generic  $\tilde{\boldsymbol{\Theta}} = (\mathbf{U}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3)$ .  $\blacksquare$

Proposition A.2 proves the case when  $N_1$  is even. If  $N_1$  is odd, we may simply consider  $N_1 - 1$  to make the largest dimension even. The result is stated in the following corollary.

**Corollary A.3.** *If  $m = 3$ ,  $N_1 \geq N_2 \geq N_3 \geq 3$ , and  $r \leq r_3$  for  $r_3$  in (4.3), then the Jacobian matrix  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank for generic  $\tilde{\Theta} = (\mathbf{U}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3)$ .*

*Proof.* When  $N_1$  is even, it is the result of Proposition A.2.

When  $N_1$  is odd, we consider the dimension  $(\tilde{N}_1, \tilde{N}_2, \tilde{N}_3)$ , where  $\tilde{N}_1, \tilde{N}_2, \tilde{N}_3$  are largest integers such that  $\tilde{N}_1$  is even,  $\tilde{N}_1 \geq \tilde{N}_2 \geq \tilde{N}_3$ , and  $N_i \geq \tilde{N}_i$  for  $i = 1, 2, 3$ . Let  $\mathbf{V}_1 := \mathbf{U}_1(1 : \tilde{N}_1, :)$ ,  $\tilde{\mathbf{V}}_2 := \tilde{\mathbf{U}}_2(1 : \tilde{N}_2 - 1, :)$ ,  $\tilde{\mathbf{V}}_3 := \tilde{\mathbf{U}}_3(1 : \tilde{N}_3 - 1, :)$ . By Proposition A.2, the matrix  $\mathbf{J}_{\tilde{\mathbf{h}}}(\mathbf{V}_1, \tilde{\mathbf{V}}_2, \tilde{\mathbf{V}}_3)$  has full column rank for generic  $\mathbf{V}_1, \tilde{\mathbf{V}}_2, \tilde{\mathbf{V}}_3$ . The matrix  $\mathbf{J}_{\tilde{\mathbf{h}}}(\mathbf{V}_1, \tilde{\mathbf{V}}_2, \tilde{\mathbf{V}}_3)$  only consists of some rows in  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$ . Thus,  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank if  $\mathbf{J}_{\tilde{\mathbf{h}}}(\mathbf{V}_1, \tilde{\mathbf{V}}_2, \tilde{\mathbf{V}}_3)$  has full column rank. It proves  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank for generic  $\tilde{\Theta} = (\mathbf{U}_1, \tilde{\mathbf{U}}_2, \tilde{\mathbf{U}}_3)$ .  $\blacksquare$

We have shown the local convexity for order-3 tensors. Finally, we prove the higher order cases by induction on the order  $m$  with base case  $m = 3$ .

**Theorem A.4.** *If  $r \leq r_m$  for  $r_m$  in (4.4), then the Jacobian matrix  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank for generic  $\tilde{\Theta} = (\mathbf{U}_1, \tilde{\mathbf{U}}_2, \dots, \tilde{\mathbf{U}}_m)$ .*

*Proof.* We will prove the result by induction on the order  $m$ . Proposition A.2 proves the conclusion for  $m = 3$ . Suppose that the result holds for the order  $m - 1$ , then we show it also holds for  $m$ . Similar to the proof of Theorem A.2, it suffices to prove there exists some  $\tilde{\Theta}$  such that the Jacobian is has full column rank [14].

Denote  $L := \min\{r_{m-1}, \lfloor \frac{n_2 n_3 \dots N_m}{N_1 + \dots + N_{m-m+1}} \rfloor\}$  and  $\mathbf{Q}_1 = \odot_{i=2}^m [\mathbf{1}^T; \tilde{\mathbf{U}}_i(:, 1 : L)]$ ,

$$\mathbf{Q}_k = \left[ \otimes_{i=2}^{k-1} [\mathbf{1}^T; \tilde{\mathbf{U}}_i(:, j)] \otimes \tilde{\mathbf{I}}_{n_k} \otimes_{i=k+1}^m [\mathbf{1}^T; \tilde{\mathbf{U}}_i(:, j)] \right]_{j=1}^L, \quad k = 2, \dots, m.$$

Let  $\mathbf{U}_1(:, i) = \mathbf{e}_j$  for  $(j - 1)L + 1 \leq i \leq jL$ , then (after omitting zero columns) it holds that

$$\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})(1 : \Pi_{i=2}^m N_i, :) = \left[ \odot_{i=2}^m [\mathbf{1}^T; \tilde{\mathbf{U}}_i(:, L + 1 : r)], \mathbf{Q}_1, \dots, \mathbf{Q}_m \right]. \quad \blacksquare$$

Let  $\mathbf{S} := [\mathbf{Q}_1, \dots, \mathbf{Q}_m]$ . The matrix  $\mathbf{S}$  has the same column space as  $\mathbf{J}_{\tilde{\mathbf{h}}}([\mathbf{1}^T; \tilde{\mathbf{U}}_2(:, 1 : L)], \dots, \tilde{\mathbf{U}}_m(:, 1 : L))$ . By the induction assumption, we know  $\mathbf{S}$  generically has full column rank since  $L \leq r_{m-1}$ . The number of columns of  $\mathbf{S}$  is  $c = (N_2 + \dots + N_m - m + 2)L$ , then

$$\begin{aligned} \Pi_{i=2}^m N_i - c &\geq \Pi_{i=2}^m N_i - (N_2 + \dots + N_m - m + 2) \lfloor \frac{\Pi_{i=2}^m N_i}{N_1 + \dots + N_{m-m+1}} \rfloor \\ &\geq \Pi_{i=2}^m N_i - \Pi_{i=2}^m N_i \frac{N_2 + \dots + N_m - m + 2}{N_1 + \dots + N_{m-m+1}} \\ &= \Pi_{i=2}^m N_i \left( 1 - \frac{N_2 + \dots + N_m - m + 2}{N_1 + \dots + N_{m-m+1}} \right) \\ &= \Pi_{i=2}^m N_i \frac{N_1 - 1}{N_1 + \dots + N_{m-m+1}} \\ &\geq (N_1 - 1) \lfloor \frac{\Pi_{i=2}^m N_i}{N_1 + \dots + N_{m-m+1}} \rfloor \\ &= r - L. \end{aligned}$$

Thus, we could choose columns  $\{\tilde{\mathbf{U}}_i(:, L+1:r)\}_{i=1}^m$  such that  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})(1:\Pi_{i=2}^m N_i, :)$  is linearly independent. It proves that  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})(1:\Pi_{i=2}^m N_i, :)$  has full column rank generically. The same proof can be applied to  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})((j-1)\Pi_{i=2}^m N_i + 1:j\Pi_{i=2}^m N_i, :), j = 1, \dots, N_1$ . Therefore,  $\mathbf{J}_{\tilde{\mathbf{h}}}(\tilde{\Theta})$  has full column rank for generic  $\tilde{\Theta} = (\mathbf{U}_1, \tilde{\mathbf{U}}_2, \dots, \tilde{\mathbf{U}}_m)$ .

Now, we are ready to prove Theorem 4.1.

*Proof.* Under the assumption of Theorem 4.1, it holds that  $\tilde{f}(\tilde{\Theta}^*) = 0$ . Therefore, Lemma A.1 and Theorem A.4 imply that the Hessian  $\nabla^2 \tilde{f}(\tilde{\Theta}^*)$  is positive definite for generic  $\tilde{\Theta}^*$ . ■

## REFERENCES

- [1] A. ABDELFAH, H. ANZT, E. G. BOMAN, E. CARSON, T. COJEAN, J. DONGARRA, A. FOX, M. GATES, N. J. HIGHAM, X. S. LI, ET AL., *A survey of numerical linear algebra methods utilizing mixed-precision arithmetic*, The International Journal of High Performance Computing Applications, 35 (2021), pp. 344–369.
- [2] A. ANANDKUMAR, R. GE, D. HSU, S. M. KAKADE, AND M. TELGARSKY, *Tensor decompositions for learning latent variable models*, Journal of machine learning research, 15 (2014), pp. 2773–2832.
- [3] C. BATTAGLINO, G. BALLARD, AND T. G. KOLDA, *A practical randomized CP tensor decomposition*, SIAM Journal on Matrix Analysis and Applications, 39 (2018), pp. 876–901.
- [4] J. BERNSTEIN, Y.-X. WANG, K. AZIZZADENESHELI, AND A. ANANDKUMAR, *signSGD: Compressed optimisation for non-convex problems*, in International Conference on Machine Learning, PMLR, 2018, pp. 560–569.
- [5] A. BEUTEL, P. P. TALUKDAR, A. KUMAR, C. FALOUTSOS, E. E. PAPALEXAKIS, AND E. P. XING, *Flexifact: Scalable flexible factorization of coupled tensors on hadoop*, in Proceedings of the 2014 SIAM international conference on data mining, SIAM, 2014, pp. 109–117.
- [6] D. BIGONI, A. P. ENGSIG-KARUP, AND Y. M. MARZOUK, *Spectral tensor-train decomposition*, SIAM Journal on Scientific Computing, 38 (2016), pp. A2405–A2439.
- [7] L. BOTTOU, F. E. CURTIS, AND J. NOCEDAL, *Optimization methods for large-scale machine learning*, Siam Review, 60 (2018), pp. 223–311.
- [8] R. BRO, *PARAFAC: tutorial and applications*, Chemometrics and intelligent laboratory systems, 38 (1997), pp. 149–171.
- [9] A. BUTTARI, J. DONGARRA, J. LANGOU, J. LANGOU, P. LUSZCZEK, AND J. KURZAK, *Mixed precision iterative refinement techniques for the solution of dense linear systems*, The International Journal of High Performance Computing Applications, 21 (2007), pp. 457–466.
- [10] E. CARSON AND N. J. HIGHAM, *Accelerating the solution of linear systems by iterative refinement in three precisions*, SIAM Journal on Scientific Computing, 40 (2018), pp. A817–A847.
- [11] E. CARSON, N. J. HIGHAM, AND S. PRANESH, *Three-precision GMRES-based iterative refinement for least squares problems*, SIAM Journal on Scientific Computing, 42 (2020), pp. A4063–A4083.
- [12] E. CARSON AND N. KHAN, *Mixed precision iterative refinement with sparse approximate inverse preconditioning*, arXiv preprint arXiv:2202.10204, (2022).
- [13] P. COMON, X. LUCIANI, AND A. L. DE ALMEIDA, *Tensor decompositions, alternating least squares and other tales*, Journal of Chemometrics: A Journal of the Chemometrics Society, 23 (2009), pp. 393–405.
- [14] D. COX, J. LITTLE, AND D. OSHEA, *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*, Springer Science & Business Media, 2013.
- [15] L. DE LATHAUWER, B. DE MOOR, AND J. VANDEWALLE, *A multilinear singular value decomposition*, SIAM journal on Matrix Analysis and Applications, 21 (2000), pp. 1253–1278.
- [16] C. DE SA, M. FELDMAN, C. RÉ, AND K. OLUKOTUN, *Understanding and optimizing asynchronous low-precision stochastic gradient descent*, in Proceedings of the 44th annual international symposium on computer architecture, 2017, pp. 561–574.
- [17] C. DE SA, M. LESZCZYNSKI, J. ZHANG, A. MARZOEV, C. R. ABERGER, K. OLUKOTUN, AND C. RÉ, *High-accuracy low-precision training*, arXiv preprint arXiv:1803.03383, (2018).

- [18] S. DOLGOV, B. N. KHOROMSKIJ, A. LITVINENKO, AND H. G. MATTHIES, *Polynomial chaos expansion of random coefficients and the solution of stochastic partial differential equations in the tensor train format*, SIAM/ASA Journal on Uncertainty Quantification, 3 (2015), pp. 1109–1135.
- [19] I. DOMANOV AND L. DE LATHAUWER, *Canonical polyadic decomposition of third-order tensors: Relaxed uniqueness conditions and algebraic algorithm*, Linear Algebra and its Applications, 513 (2017), pp. 342–375.
- [20] M. DRESSLER, J. NIE, AND Z. YANG, *Separability of hermitian tensors and psd decompositions*, Linear and Multilinear Algebra, (2021), pp. 1–28.
- [21] R. GE AND T. MA, *On the optimization landscape of tensor decompositions*, Advances in Neural Information Processing Systems, 30 (2017).
- [22] L. GRASEDYCK, *Hierarchical singular value decomposition of tensors*, SIAM journal on matrix analysis and applications, 31 (2010), pp. 2029–2054.
- [23] B. GUO, J. NIE, AND Z. YANG, *Learning diagonal gaussian mixture models and incomplete tensor decompositions*, Vietnam Journal of Mathematics, 50 (2022), pp. 421–446.
- [24] A. HAIDAR, H. BAYRAKTAR, S. TOMOV, J. DONGARRA, AND N. J. HIGHAM, *Mixed-precision iterative refinement using tensor cores on GPUs to accelerate solution of linear systems*, Proceedings of the Royal Society A, 476 (2020), p. 20200110.
- [25] C. HAWKINS, X. LIU, AND Z. ZHANG, *Towards compact neural networks via end-to-end training: A Bayesian tensor approach with automatic rank determination*, SIAM Journal on Mathematics of Data Science, 4 (2022), pp. 46–71.
- [26] C. HAWKINS AND Z. ZHANG, *Bayesian tensorized neural networks with automatic rank selection*, Neurocomputing, 453 (2021), pp. 172–180.
- [27] D. L. N. HETTIARACHCHI, V. S. P. DAVULURU, AND E. J. BALSTER, *Integer vs. floating-point processing on modern fpga technology*, in 2020 10th Annual Computing and Communication Workshop and Conference (CCWC), IEEE, 2020, pp. 0606–0612.
- [28] R. A. HORN AND C. R. JOHNSON, *Matrix analysis*, Cambridge university press, 2012.
- [29] I. HUBARA, M. COURBARIAUX, D. SOUDRY, R. EL-YANIV, AND Y. BENGIO, *Quantized neural networks: Training neural networks with low precision weights and activations*, The Journal of Machine Learning Research, 18 (2017), pp. 6869–6898.
- [30] W. HUGGINS, P. PATIL, B. MITCHELL, K. B. WHALEY, AND E. M. STOUDENMIRE, *Towards quantum machine learning with tensor networks*, Quantum Science and technology, 4 (2019), p. 024001.
- [31] N. P. JOUPPI, C. YOUNG, N. PATIL, D. PATTERSON, G. AGRAWAL, R. BAJWA, S. BATES, S. BHATIA, N. BODEN, A. BORCHERS, ET AL., *In-datacenter performance analysis of a tensor processing unit*, in Proceedings of the 44th annual international symposium on computer architecture, 2017, pp. 1–12.
- [32] A. KERR, H. WU, M. GUPTA, D. BLASIG, P. RAMINI, D. MERRILL, A. SHIVAM, P. MAJCHER, P. SPRINGER, M. HOHNERBACH, J. WANG, AND M. NICELY, *CUTLASS*, 4 2022, <https://github.com/NVIDIA/cutlass>.
- [33] Y.-D. KIM, E. PARK, S. YOO, T. CHOI, L. YANG, AND D. SHIN, *Compression of deep convolutional neural networks for fast and low power mobile applications*, arXiv preprint arXiv:1511.06530, (2015).
- [34] F. KNOLL, J. ZBONTAR, A. SRIRAM, M. J. MUCKLEY, M. BRUNO, A. DEFAZIO, M. PARENTE, K. J. GERAS, J. KATSNELSON, H. CHANDARANA, ET AL., *fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning*, Radiology. Artificial intelligence, 2 (2020).
- [35] T. G. KOLDA AND B. W. BADER, *Tensor decompositions and applications*, SIAM review, 51 (2009), pp. 455–500.
- [36] T. G. KOLDA AND D. HONG, *Stochastic gradients for large-scale tensor decomposition*, SIAM Journal on Mathematics of Data Science, 2 (2020), pp. 1066–1095.
- [37] J. M. LANDSBERG, *Tensors: geometry and applications*, Representation theory, 381 (2012), p. 3.
- [38] S. A. NENE, S. K. NAYAR, H. MURASE, ET AL., *Columbia object image library (coil-100)*, (1996).
- [39] J. NIE, *Generating polynomials and symmetric tensor decompositions*, Foundations of Computational Mathematics, 17 (2017), pp. 423–465.
- [40] J. NIE AND Z. YANG, *Hermitian tensor decompositions*, SIAM Journal on Matrix Analysis and Applications, 41 (2020), pp. 1115–1144.
- [41] J. NIE, Z. YANG, AND X. ZHANG, *A complete semidefinite algorithm for detecting copositive matrices*



- and tensors, *SIAM Journal on Optimization*, 28 (2018), pp. 2902–2921.
- [42] A. NOVIKOV, D. PODOPRIKHIN, A. OSOKIN, AND D. P. VETROV, *Tensorizing neural networks*, *Advances in neural information processing systems*, 28 (2015).
- [43] R. OKUTA, Y. UNNO, D. NISHINO, S. HIDO, AND C. LOOMIS, *CuPy: A NumPy-compatible library for NVIDIA GPU calculations*, in *Proc. Workshop on Machine Learning Systems in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [44] R. OLIVARES-AMAYA, M. WATSON, R. EDGAR, L. VOGT, Y. SHAO, AND A. ASPURU-GUZIK, *Accelerating correlated quantum chemistry calculations using graphical processing units and a mixed precision matrix multiplication library*, *Journal of chemical theory and computation*, 6 (2010), pp. 135–144.
- [45] R. ORÚS, *Tensor networks for complex quantum systems*, *Nature Reviews Physics*, 1 (2019), pp. 538–550.
- [46] I. V. OSELEDETS, *Tensor-train decomposition*, *SIAM Journal on Scientific Computing*, 33 (2011), pp. 2295–2317.
- [47] L. RICHTER, L. SALLANDT, AND N. NÜSKEN, *Solving high-dimensional parabolic PDEs using the tensor train format*, in *International Conference on Machine Learning*, PMLR, 2021, pp. 8998–9009.
- [48] J. SHAN, M. R. CASU, J. CORTADELLA, L. LAVAGNO, AND M. T. LAZARESCU, *Exact and heuristic allocation of multi-kernel applications to multi-FPGA platforms*, in *Proceedings of the 56th Annual Design Automation Conference 2019*, 2019, pp. 1–6.
- [49] N. D. SIDIROPOULOS, L. DE LATHAUWER, X. FU, K. HUANG, E. E. PAPALEXAKIS, AND C. FALOUTSOS, *Tensor decomposition for signal processing and machine learning*, *IEEE Transactions on Signal Processing*, 65 (2017), pp. 3551–3582.
- [50] X. SUN, N. WANG, C.-Y. CHEN, J. NI, A. AGRAWAL, X. CUI, S. VENKATARAMANI, K. EL MAGHRAOUI, V. V. SRINIVASAN, AND K. GOPALAKRISHNAN, *Ultra-low precision 4-bit training of deep neural networks*, *Advances in Neural Information Processing Systems*, 33 (2020), pp. 1796–1807.
- [51] N. VERVLIET AND L. DE LATHAUWER, *A randomized block sampling approach to canonical polyadic decomposition of large-scale tensors*, *IEEE Journal of Selected Topics in Signal Processing*, 10 (2015), pp. 284–295.
- [52] B. YAMAN, S. WEINGÄRTNER, N. KARGAS, N. D. SIDIROPOULOS, AND M. AKÇAKAYA, *Low-rank tensor models for improved multidimensional MRI: Application to dynamic cardiac  $t_1$  mapping*, *IEEE transactions on computational imaging*, 6 (2019), pp. 194–207.
- [53] J. ZBONTAR, F. KNOLL, A. SRIRAM, T. MURRELL, Z. HUANG, M. J. MUCKLEY, A. DEFAZIO, R. STERN, P. JOHNSON, M. BRUNO, ET AL., *fastMRI: An open dataset and benchmarks for accelerated MRI*, *arXiv preprint arXiv:1811.08839*, (2018).
- [54] K. ZHANG, C. HAWKINS, X. ZHANG, C. HAO, AND Z. ZHANG, *On-FPGA training with ultra memory reduction: A low-precision tensor method*, *arXiv preprint arXiv:2104.03420*, (2021).
- [55] Z. ZHANG, T.-W. WENG, AND L. DANIEL, *Big-data tensor recovery for high-dimensional uncertainty quantification of process variations*, *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 7 (2016), pp. 687–697.
- [56] Z. ZHANG, X. YANG, I. V. OSELEDETS, G. E. KARNIADAKIS, AND L. DANIEL, *Enabling high-dimensional hierarchical uncertainty quantification by ANOVA and tensor-train decomposition*, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34 (2014), pp. 63–76.