

Parallel Numerical Algorithms

Chapter 7 – Differential Equations

Section 7.5 – Tensor Analysis

Edgar Solomonik

Department of Computer Science
University of Illinois at Urbana-Champaign

CS 554 / CSE 512

Outline

- 1 Tensor Algebra
 - Tensors
 - Tensor Transposition
 - Tensor Contractions
- 2 Tensor Decompositions
 - CP Decomposition
 - Tucker Decomposition
 - Tensor Train Decomposition
- 3 Fast Algorithms
 - Strassen's Algorithm
 - Bilinear Algorithms

Tensors

A *tensor* $\mathbf{T} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ has

- *Order* d (i.e. d *modes / indices*)
- *Dimensions* n_1 -by- \dots -by- n_d
- *Elements* $t_{i_1 \dots i_d} = t_{\mathbf{i}}$ where $\mathbf{i} \in \bigotimes_{i=1}^d \{1, \dots, n_i\}$

Order d tensors represent d -dimensional arrays

- ($d \geq 3$)-dimensional arrays are prevalent in scientific computing
 - Regular grids, collections of matrices, multilinear operators
 - Experimental data, visual/graphic data
- Tensors analysis is the expression and study of numerical methods using tensor representations

Reshaping Tensors

When using tensors, it is often necessary to transition between high-order and low-order representations of the same object

- Recall for a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ its *unfolding* is given by

$$\mathbf{v} = \text{vec}(\mathbf{A}), \Rightarrow \mathbf{v} \in \mathbb{R}^{mn}, v_{i+jm} = a_{ij}$$

- A tensor $\mathbf{T} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ can be fully unfolded the same way

$$\mathbf{v} = \text{vec}(\mathbf{T}), \Rightarrow \mathbf{v} \in \mathbb{R}^{n_1 \dots n_d}, v_{i_1+i_2n_1+i_3n_1n_2+\dots} = t_{i_1i_2i_3\dots}$$

- Often we also want to *fold* tensors into higher-order ones
- Generally, we can *reshape* (fold or unfold) any tensor

$$\mathbf{U} = o_{n_1 \times \dots \times n_d}(\mathbf{V}) \quad \Rightarrow \quad \mathbf{U} \in \mathbb{R}^{n_1 \times \dots \times n_d}, \quad \text{vec}(\mathbf{U}) = \text{vec}(\mathbf{V})$$

Tensor Transposition

For tensors of order ≥ 3 , there is more than one way to transpose modes

- A *tensor transposition* is defined by a permutation p containing elements $\{1, \dots, d\}$

$$\mathbf{Y} = \mathbf{X}^{\langle p \rangle} \quad \Rightarrow \quad y_{i_{p_1}, \dots, i_{p_d}} = x_{i_1, \dots, i_d}$$

- In this notation, a transposition of matrix \mathbf{A} is defined as

$$\mathbf{A}^T = \mathbf{A}^{\langle [2,1] \rangle}$$

- Tensor transposition is a convenient primitive for manipulating multidimensional arrays and mapping tensor computations to linear algebra
- In tensor derivations, indices are often carried through to avoid transpositions

Tensor Symmetry

We say a tensor is *symmetric* if $\forall j, k \in \{1, \dots, d\}$

$$t_{i_1 \dots i_j \dots i_k \dots i_d} = t_{i_1 \dots i_k \dots i_j \dots i_d} \quad \text{or equivalently} \quad \mathbf{T} = \mathbf{T}^{\langle [1, \dots, j, \dots, k, \dots, d] \rangle}$$

A tensor is *antisymmetric* (skew-symmetric) if $\forall j, k \in \{1, \dots, d\}$

$$t_{i_1 \dots i_j \dots i_k \dots i_d} = (-1) t_{i_1 \dots i_k \dots i_j \dots i_d}$$

A tensor is *partially-symmetric* if such index interchanges are restricted to be within subsets of $\{1, \dots, d\}$, e.g.

$$t_{ijkl} = t_{jikl} = t_{jilk} = t_{ijlk}$$

Tensor Products and Kronecker Products

Tensor products can be defined with respect to maps

$$f : V_f \rightarrow W_f \text{ and } g : V_g \rightarrow W_g$$

$$h = f \times g \quad \Rightarrow \quad g : (V_f \times V_g) \rightarrow (W_f \times W_g), \quad h(x, y) = f(x)g(y)$$

Tensors can be used to represent multilinear maps and have a corresponding definition for a tensor product

$$\mathbf{T} = \mathbf{X} \times \mathbf{Y} \quad \Rightarrow \quad t_{i_1, \dots, i_m, j_1, \dots, j_n} = x_{i_1, \dots, i_m} y_{j_1, \dots, j_n}$$

The *Kronecker product* between two matrices $\mathbf{A} \in \mathbb{R}^{m_1 \times m_2}$, $\mathbf{B} \in \mathbb{R}^{n_1 \times n_2}$

$$\mathbf{C} = \mathbf{A} \otimes \mathbf{B} \quad \Rightarrow \quad c_{i_2+i_1m_2, j_2+j_1n_2} = a_{i_1j_1} b_{i_2j_2}$$

corresponds to transposing and reshaping the tensor product

$$\mathbf{A} \otimes \mathbf{B} = o_{m_1n_1, m_2n_2}((\mathbf{A} \times \mathbf{B})^{\langle [3,1,4,2] \rangle})$$

Tensor Partial Sum

- \mathbf{Y} of order $d - r$ is a *partial sum* of \mathbf{X} of order d if for some \mathbf{q} containing r elements of $\{1, \dots, d\}$

$$\mathbf{Y} = \sum_{\mathbf{q}} (\mathbf{X}) \quad \Rightarrow \quad \bar{\mathbf{X}} = \mathbf{X} \langle [q_1, \dots, q_r, \dots] \rangle,$$
$$y_{i_1 \dots i_{d-r}} = \sum_{j_1} \cdots \sum_{j_r} \bar{x}_{j_1, \dots, j_r, i_1, \dots, i_{d-r}}$$

- Partial summations provide a powerful primitive operation when coupled with transposition and reshape

Tensor Trace

- Z of order $d - 2r$ is a *trace* of X of order $d \geq r$ if for some p, q each containing a different set of r elements of $\{1, \dots, d\}$

$$Y = \text{trace}_{p,q}(X) \quad \Rightarrow \quad \bar{X} = X \langle [p_1, \dots, p_r, q_1, \dots, q_r, \dots] \rangle,$$

$$y_{i_1 \dots i_{d-2r}} = \sum_{j_1} \cdots \sum_{j_r} \bar{x}_{j_1, \dots, j_r, j_1, \dots, j_r, i_1, \dots, i_{d-2r}}$$

- The trace of a matrix A in this notation is

$$\text{trace}(A) = \text{trace}_{[0],[1]}(A) = \sum_i a_{ii}$$

Tensor Contraction

Tensor contraction is a transpose of a trace of a tensor product

$$C = \left[\text{trace}_{p,q}(A \times B) \right]^{\langle r \rangle} \quad \text{for some } p, q, r$$

- Examples in linear algebra include: vector inner and outer products, matrix–vector product, matrix–matrix product
- The *contracted* modes of A appear in p and of B in q , while *uncontracted* modes appear in r
- Matrix multiplication would be given by $p = [2]$, $q = [3]$, $r = [1, 4]$

Tensor Times Matrix

Tensor times matrix (TTM) is one of the most common tensor contractions involving tensors of order ≥ 3

- Given an order 3 tensor \mathbf{T} and matrix \mathbf{V} , TTM computes order 3 tensor \mathbf{W} , generalizes naturally to higher-order \mathbf{T}
- TTM can contract one of three modes of \mathbf{T}

$$\mathbf{W} = \left[\text{trace}_{[3],[4]}(\mathbf{T} \times \mathbf{V}) \right]^{\langle [1,2,5] \rangle} \quad \text{or} \quad \mathbf{W} = \left[\text{trace}_{[2],[4]}(\mathbf{T} \times \mathbf{V}) \right]^{\langle [1,3,5] \rangle}$$
$$\text{or} \quad \mathbf{W} = \left[\text{trace}_{[1],[4]}(\mathbf{T} \times \mathbf{V}) \right]^{\langle [2,3,5] \rangle}$$

- In the first case, we have

$$w_{ijk} = \sum_l t_{ijl} v_{lk}$$

Tensor Contraction Diagrams


Consider the tensor contraction

$$w_{ijk} = \sum_{lm} t_{ijlm} v_{mkl}$$

which we can also write in tensor notation

$$\mathbf{W} = \left[\text{trace}_{[3,4],[7,5]} (\mathbf{T} \times \mathbf{V}) \right]^{\langle [1,2,6] \rangle}$$

or in the following diagrammatic form



CP decomposition

The SVD corresponds to a sum of outer products of the form

$$\mathbf{A} = \sum_{k=1}^r \sigma_k \mathbf{u}_k \mathbf{v}_k^T$$

so it is natural to seek to approximate a tensor as

$$\mathbf{T} = \sum_{k=1}^r \sigma_k \mathbf{u}_k^{(1)} \times \cdots \times \mathbf{u}_k^{(d)}$$

where each $\mathbf{u}_k^{(i)}$ is orthogonal to any other $\mathbf{u}_{k'}^{(i)}$, yielding the *canonical polyadic (CP) decomposition* of \mathbf{T}


- r is referred to as the *canonical rank* of the tensor

Computing the CP decomposition

- Computing the canonical rank is *NP hard*
- Approximation by CP decomposition is *ill-posed*
- Regularization (imposing bounds on the norm of the factor matrices) make the optimization problem feasible
- *Alternating least squares (ALS)* commonly used for computation
 - Optimizes for one factor matrix at a time
 - Least squares problem for each matrix
- Alternatives include coordinate and gradient descent methods, much like in numerical optimization for matrix completion

Tucker Decomposition

The *Tucker decomposition* introduces an order d *core tensor* into the CP decomposition



$$t_{i_1 \dots i_d} = \sum_{k_1 \dots k_d} s_{k_1 \dots k_d} w_{i_1 k_1}^{(1)} \cdots w_{i_d k_d}^{(d)}$$

where the columns of each $\mathbf{W}^{(i)}$ are orthonormal

- Unlike CP decomposition (given by ‘diagonal’ tensor S), each index appears in no more than two tensors
- Tucker decomposition is not *low-order* since the order of T matches that of S

Computing the Tucker Decomposition

The *SVD (HOSVD)* can refer to the Tucker decomposition or the following basic method for its computation

- Compute the left singular vectors $\mathbf{W}^{(i)}$ of all d single-mode unfoldings of \mathbf{T}

$$\mathbf{T}_{(i)} = o_{n_i \times n_1 \cdots n_{i-1} n_{i+1} \cdots n_d} (\mathbf{T}^{\langle [i, 1, \dots, i-1, i+1, \dots, d] \rangle})$$


- Compute the core tensor by projecting \mathbf{A} onto these singular vectors along each mode

$$\mathbf{S} = \left[\begin{array}{c} \text{trace} \\ [1, \dots, d], [d+2, d+4, \dots, 2d] \end{array} (\mathbf{T} \times \mathbf{W}^{(1)} \times \dots \times \mathbf{W}^{(d)}) \right]^{\langle [d+1, d+3, \dots, 2d-1] \rangle}$$

- The HOSVD works well when the core tensor \mathbf{S} can be shown to have decaying entries

Tensor Train Decomposition

The *tensor train* decomposition has the following diagrammatic representation



- The tensor train is a chain of contracted order 3 tensors with the two ends having order 2
- Elements of T are given by matrix product chain

$$t_{i_1 \dots i_d} = \mathbf{u}^{(i_1)} \mathbf{U}^{(i_2)} \dots \mathbf{U}^{(i_{d-1})} \mathbf{u}^{(i_d)}$$

- Has been used for decades in physics, known as the *matrix product states (MPS)* representation

Tensor Train Properties

- The *tensor train (TT) ranks* are given by the dimensions of the auxiliary modes in the factorization
- The *tensor train (TT) rank* is the maximum number of columns in any matrix $\mathbf{U}^{(i_j)}$
- The tensor train rank is a matrix rank, i.e. it corresponds to a low-rank decomposition of a matrix (given by contracting two parts of the tensor train)
- Summation and products of tensor trains can be readily computed

Quantized Tensor Train

- The *quantized tensor train (QTT)* corresponds to the application of TT to a tensor that is reshaped to be higher-order (e.g. each resulting mode dimension is constant)
- For some classes of matrices QTT analytic decompositions are known
- Toeplitz matrices have constant TT rank
- 3D Poisson operator has constant TT rank, more generally for FEM methods with simple mass and stiffness matrices

Computing the Tensor Train Decomposition

- Product of matrix with constant QTT rank and vector of dimension n has cost $\Theta(n \log n)$
- Given general vector, can compute TT decomposition by hierarchical SVD
- Faster algorithms leveraging low-rank of SVD are possible
- Can interpolate order d tensor with dimensions equal to n and TT rank r with cost $O(dnr^2)$
- Can efficiently obtain *cross approximation*, i.e. a lower-rank approximation to an existing TT approximation
- For order d tensor with n -dimensions and TT rank r , cross approximation cost is $O(dnr^3)$

Strassen's Algorithm

$$\text{Strassen's algorithm } \begin{bmatrix} C_{11} & C_{12} \\ C_{21} & C_{22} \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \cdot \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix}$$

$$M_1 = (A_{11} + A_{22}) \cdot (B_{11} + B_{22})$$

$$C_{11} = M_1 + M_4 - M_5 + M_7$$

$$M_2 = (A_{21} + A_{22}) \cdot B_{11}$$

$$C_{21} = M_2 + M_4$$

$$M_3 = A_{11} \cdot (B_{12} - B_{22})$$

$$C_{12} = M_3 + M_5$$

$$M_4 = A_{22} \cdot (B_{21} - B_{11})$$

$$C_{22} = M_1 - M_2 + M_3 + M_6$$

$$M_5 = (A_{11} + A_{12}) \cdot B_{22}$$

$$M_6 = (A_{21} - A_{11}) \cdot (B_{11} + B_{12})$$

$$M_7 = (A_{12} - A_{22}) \cdot (B_{21} + B_{22})$$

Minimize products \Rightarrow minimize number of recursive calls

$$T(n) = 7T(n/2) + O(n^2) = O(7^{\log_2 n}) = O(n^{\log_2 7})$$

For convolution, DFT matrix reduces from naive $O(n^2)$ products to $O(n)$, both of these are **bilinear algorithms**

Bilinear Algorithms

Definition (Bilinear algorithms (V. Pan, 1984))

A bilinear algorithm $\Lambda = (\mathbf{F}^{(A)}, \mathbf{F}^{(B)}, \mathbf{F}^{(C)})$ computes

$$\mathbf{c} = \mathbf{F}^{(C)}[(\mathbf{F}^{(A)\top} \mathbf{a}) \odot (\mathbf{F}^{(B)\top} \mathbf{b})],$$

where \mathbf{a} and \mathbf{b} are inputs and \odot is the Hadamard (pointwise) product.

$$\begin{bmatrix} \mathbf{c} \end{bmatrix} = \begin{bmatrix} \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \end{bmatrix} \left[\left(\begin{bmatrix} \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \end{bmatrix}^\top \begin{bmatrix} \mathbf{a} \end{bmatrix} \right) \odot \left(\begin{bmatrix} \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times \end{bmatrix}^\top \begin{bmatrix} \mathbf{b} \end{bmatrix} \right) \right]$$

Bilinear Algorithms as Tensor Factorizations

A bilinear algorithm corresponds to a CP tensor decomposition

$$\begin{aligned}c_i &= \sum_{r=1}^r f_{ir}^{(C)} \left(\sum_j f_{jr}^{(A)} \mathbf{a}_j \right) \left(\sum_k f_{kr}^{(B)} b_k \right) \\ &= \sum_j \sum_k \left(\sum_{r=1}^r f_{ir}^{(C)} f_{jr}^{(A)} f_{kr}^{(B)} \right) a_j b_k \\ &= \sum_j \sum_k t_{ijk} a_j b_k \quad \text{where} \quad t_{ijk} = \sum_{r=1}^r f_{ir}^{(C)} f_{jr}^{(A)} f_{kr}^{(B)}\end{aligned}$$

For multiplication of $n \times n$ matrices,

- \mathbf{T} is $n^2 \times n^2 \times n^2$
- Classical algorithm has rank $r = n^3$
- Strassen's algorithm has rank $r \approx n^{\log_2(7)}$

References

- Golub, Gene H., and Charles F. Van Loan. Matrix computations. Vol. 3. *JHU Press*, 2012.
- Kolda, Tamara G., and Brett W. Bader. Tensor decompositions and applications. *SIAM review* 51.3 (2009): 455-500.
- Khoromskij, Boris N. Tensors-structured numerical methods in scientific computing: Survey on recent advances. *Chemometrics and Intelligent Laboratory Systems* 110.1 (2012): 1-19.
- Grasedyck, Lars, Daniel Kressner, and Christine Tobler. A literature survey of low-rank tensor approximation techniques. *GAMM-Mitteilungen* 36.1 (2013): 53-78.

References

- Kazeev, Vladimir A., and Boris N. Khoromskij. Low-rank explicit QTT representation of the Laplace operator and its inverse. *SIAM Journal on Matrix Analysis and Applications* 33.3 (2012): 742-758.
- Khoromskaia, Venera, Boris Khoromskij, and Reinhold Schneider. QTT representation of the Hartree and exchange operators in electronic structure calculations. *Computational Methods in Applied Mathematics Comput. Methods Appl. Math.* 11.3 (2011): 327-341.
- Khoromskij, Boris N., and Ivan V. Oseledets. QTT approximation of elliptic solution operators in higher dimensions. *Russian Journal of Numerical Analysis and Mathematical Modelling* 26.3 (2011): 303-322.

References

- De Silva, Vin, and Lek-Heng Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM Journal on Matrix Analysis and Applications* 30.3 (2008): 1084-1127.
- Pan, Victor. How can we speed up matrix multiplication? *SIAM review* 26.3 (1984): 393-415.
- Karlsson, Lars, Daniel Kressner, and André Uschmajew. Parallel algorithms for tensor completion in the CP format. *Parallel Computing* 57 (2016): 222-234.