

Nonnegative Tensor Completion

- Let $\mathcal{X} \in \mathbb{R}_+^{I \times J \times K}$ be an incomplete tensor, and $\Omega \subseteq \{1 \dots I\} \times \{1 \dots J\} \times \{1 \dots K\}$ be the set of indices of its known entries [1].
- Also, let $\mathcal{M} \in \mathbb{R}^{I \times J \times K}$ with

$$\mathcal{M}(i, j, k) = \begin{cases} 1, & \text{if } (i, j, k) \in \Omega, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

- We consider the Nonnegative Tensor Completion (NTC) problem

$$\min_{\mathbf{A} \geq 0, \mathbf{B} \geq 0, \mathbf{C} \geq 0} f_{\Omega}(\mathbf{A}, \mathbf{B}, \mathbf{C}) + \frac{\lambda}{2} \|\mathbf{A}\|_F^2 + \frac{\lambda}{2} \|\mathbf{B}\|_F^2 + \frac{\lambda}{2} \|\mathbf{C}\|_F^2, \quad (2)$$

where $\mathbf{A} = [\mathbf{a}_1 \dots \mathbf{a}_I] \in \mathbb{R}_+^{I \times R}$, $\mathbf{B} = [\mathbf{b}_1 \dots \mathbf{b}_J] \in \mathbb{R}_+^{J \times R}$, $\mathbf{C} = [\mathbf{c}_1 \dots \mathbf{c}_K] \in \mathbb{R}_+^{K \times R}$, and

$$f_{\Omega}(\mathbf{A}, \mathbf{B}, \mathbf{C}) = \frac{1}{2} \|\mathcal{M} \circledast (\mathcal{X} - [\mathbf{A}, \mathbf{B}, \mathbf{C}])\|_F^2, \quad (3)$$

with

$$[\mathbf{A}, \mathbf{B}, \mathbf{C}] = \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r. \quad (4)$$

Alternating optimization framework

- We can derive matrix-based equivalent expressions of f_{Ω} as

$$\begin{aligned} f_{\Omega}(\mathbf{A}, \mathbf{B}, \mathbf{C}) &= \frac{1}{2} \|\mathbf{M}_{\mathbf{A}} \circledast (\mathbf{X}_{\mathbf{A}} - \mathbf{A}(\mathbf{C} \circledast \mathbf{B})^T)\|_F^2 = \frac{1}{2} \|\mathbf{M}_{\mathbf{B}} \circledast (\mathbf{X}_{\mathbf{B}} - \mathbf{B}(\mathbf{C} \circledast \mathbf{A})^T)\|_F^2 \\ &= \frac{1}{2} \|\mathbf{M}_{\mathbf{C}} \circledast (\mathbf{X}_{\mathbf{C}} - \mathbf{C}(\mathbf{B} \circledast \mathbf{A})^T)\|_F^2, \end{aligned} \quad (5)$$

where $\mathbf{M}_{\mathbf{A}}$, $\mathbf{M}_{\mathbf{B}}$, $\mathbf{M}_{\mathbf{C}}$, and $\mathbf{X}_{\mathbf{A}}$, $\mathbf{X}_{\mathbf{B}}$, $\mathbf{X}_{\mathbf{C}}$ are the matrix unfoldings of \mathcal{M} and \mathcal{X} , respectively.

- Solving (2) for \mathbf{A} , \mathbf{B} , \mathbf{C} is a non-convex problem.

- Alternating optimization (AO):

- Initialize $\mathbf{A}_0, \mathbf{B}_0, \mathbf{C}_0, l = 0$.

$$1 \quad \mathbf{A}_{l+1} = \underset{\mathbf{A} \geq 0}{\operatorname{argmin}} f_{\Omega}(\mathbf{A}) := \frac{1}{2} \|\mathbf{M}_{\mathbf{A}} \circledast (\mathbf{X}_{\mathbf{A}} - \mathbf{A}(\mathbf{C}_l \circledast \mathbf{B}_l)^T)\|_F^2 + \frac{\lambda}{2} \|\mathbf{A}\|_F^2.$$

$$2 \quad \mathbf{B}_{l+1} = \underset{\mathbf{B} \geq 0}{\operatorname{argmin}} f_{\Omega}(\mathbf{B}) := \frac{1}{2} \|\mathbf{M}_{\mathbf{B}} \circledast (\mathbf{X}_{\mathbf{B}} - \mathbf{B}(\mathbf{C}_l \circledast \mathbf{A}_{l+1})^T)\|_F^2 + \frac{\lambda}{2} \|\mathbf{B}\|_F^2.$$

$$3 \quad \mathbf{C}_{l+1} = \underset{\mathbf{C} \geq 0}{\operatorname{argmin}} f_{\Omega}(\mathbf{C}) := \frac{1}{2} \|\mathbf{M}_{\mathbf{C}} \circledast (\mathbf{X}_{\mathbf{C}} - \mathbf{C}(\mathbf{B}_{l+1} \circledast \mathbf{A}_{l+1})^T)\|_F^2 + \frac{\lambda}{2} \|\mathbf{C}\|_F^2.$$

- Iterate 1, 2, 3 until convergence.

Nonnegative Matrix Completion

- Let $\mathbf{X} \in \mathbb{R}_+^{m \times n}$ be an incomplete matrix, and $\Omega \subseteq \{1 \dots m\} \times \{1 \dots n\}$ be the set of indices of its known entries.

- Also, let $\mathbf{A} \in \mathbb{R}_+^{m \times r}$, $\mathbf{B} \in \mathbb{R}_+^{n \times r}$, and $\mathbf{M} \in \mathbb{R}^{m \times n}$ with

$$\mathbf{M}(i, j) = \begin{cases} 1, & \text{if } (i, j) \in \Omega, \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

- We consider the Nonnegative Matrix Completion (NMC) problem

$$\min_{\mathbf{A} \geq 0} f_{\Omega}(\mathbf{A}) := \frac{1}{2} \|\mathbf{M} \circledast (\mathbf{X} - \mathbf{A}\mathbf{B}^T)\|_F^2 + \frac{\lambda}{2} \|\mathbf{A}\|_F^2. \quad (7)$$

- The gradient and the Hessian of f_{Ω} , at point \mathbf{A} , are given by

$$\nabla f_{\Omega}(\mathbf{A}) = -(\mathbf{M} \circledast \mathbf{X} - \mathbf{M} \circledast \mathbf{A}\mathbf{B}^T) \mathbf{B} + \lambda \mathbf{A}, \quad (8)$$

and

$$\nabla^2 f_{\Omega}(\mathbf{A}) = (\mathbf{B}^T \circledast \mathbf{I}) \operatorname{diag}^2(\operatorname{vec}(\mathbf{M})) (\mathbf{B} \circledast \mathbf{I}) + \lambda \mathbf{I}. \quad (9)$$

Nesterov-Type Algorithm for NMC

Algorithm 1: $\mathbf{A}_{\text{opt}} = \text{N_NMC}(\mathbf{X}, \mathbf{M}, \mathbf{B}, \mathbf{A}_*)$

Input: $\mathbf{X}, \mathbf{M} \in \mathbb{R}_+^{m \times n}$, $\mathbf{B} \in \mathbb{R}_+^{n \times r}$, $\mathbf{A}_* \in \mathbb{R}_+^{m \times r}$, λ, μ, L

$\mathbf{W} = -(\mathbf{M} \circledast \mathbf{X})\mathbf{B}$

$q = \frac{\mu + \lambda}{L + \lambda}$

$\mathbf{A}_0 = \mathbf{Y}_0 = \mathbf{A}_*$

$\alpha_0 = 1, l = 0$

while (1) **do**

$\nabla f_{\Omega}(\mathbf{Y}_l) = \mathbf{W} + (\mathbf{M} \circledast \mathbf{Y}_l \mathbf{B}^T) \mathbf{B} + \lambda \mathbf{Y}_l$

if (term_cond is TRUE) **then**

break

else

$\mathbf{A}_{l+1} = (\mathbf{Y}_l - \frac{1}{L + \lambda} \nabla f_{\Omega}(\mathbf{Y}_l))_+$

$\alpha_{l+1}^2 = (1 - \alpha_{l+1})\alpha_l^2 + q\alpha_{l+1}$

$\beta_{l+1} = \frac{\alpha_l(1 - \alpha_l)}{\alpha_{l+1}^2 + \alpha_{l+1}}$

$\mathbf{Y}_{l+1} = \mathbf{A}_{l+1} + \beta_{l+1}(\mathbf{A}_{l+1} - \mathbf{A}_l)$

$l = l + 1$

return \mathbf{A}_l .

Nesterov Based AO NTC

Algorithm 2: Nesterov-based AO NTC

Input: $\mathcal{X}, \Omega, \mathbf{A}_0 > 0, \mathbf{B}_0 > 0, \mathbf{C}_0 > 0$.

$l = 0$

while (1) **do**

$\mathbf{A}_{l+1} = \text{N_NMC}(\mathbf{X}_{\mathbf{A}}, \mathbf{M}_{\mathbf{A}}, (\mathbf{C}_l \circledast \mathbf{B}_l), \mathbf{A}_l)$

$\mathbf{B}_{l+1} = \text{N_NMC}(\mathbf{X}_{\mathbf{B}}, \mathbf{M}_{\mathbf{B}}, (\mathbf{C}_l \circledast \mathbf{A}_{l+1}), \mathbf{B}_l)$

$\mathbf{C}_{l+1} = \text{N_NMC}(\mathbf{X}_{\mathbf{C}}, \mathbf{M}_{\mathbf{C}}, (\mathbf{A}_{l+1} \circledast \mathbf{B}_{l+1}), \mathbf{C}_l)$

if (term_cond is TRUE) **then break; endif**

$l = l + 1$

return $\mathbf{A}_l, \mathbf{B}_l, \mathbf{C}_l$.

Computation of $\mathbf{W}_{\mathbf{A}}$ and $\mathbf{Z}_{\mathbf{A}}$

$$\mathbf{W}_{\mathbf{A}} = (\mathbf{M}_{\mathbf{A}} \circledast \mathbf{X}_{\mathbf{A}})(\mathbf{C} \circledast \mathbf{B}), \quad \mathbf{Z}_{\mathbf{A}} = (\mathbf{M}_{\mathbf{A}} \circledast \mathbf{A}(\mathbf{C} \circledast \mathbf{B})^T)(\mathbf{C} \circledast \mathbf{B}). \quad (10)$$

- The i -th row of $\mathbf{W}_{\mathbf{A}}$, for $i = 1, \dots, I$, is computed as

$$\mathbf{W}_{\mathbf{A}}(i, :) = (\mathbf{M}_{\mathbf{A}}(i, :) \circledast \mathbf{X}_{\mathbf{A}}(i, :))(\mathbf{C} \circledast \mathbf{B}). \quad (11)$$

- The computation involves the multiplication of a $(1 \times JK)$ row vector and a $(JK \times R)$ matrix.

- In order to reduce the computational complexity, we must exploit the sparsity of \mathcal{X} .

- Let nnz_i be the number of known entries in the i -th horizontal slice of \mathcal{X} . Also, let these known entries have indices $(i, j_q, k_q) \in \Omega$, for $q = 1, \dots, nnz_i$.

- The computation of the i -th row of $\mathbf{W}_{\mathbf{A}}$ reduces to

$$\mathbf{W}_{\mathbf{A}}(i, :) = \sum_{q=1}^{nnz_i} \mathcal{X}(i, j_q, k_q) \mathbf{C}(k_q, :) \circledast \mathbf{B}(j_q, :). \quad (12)$$

Efficient computation of $\mathbf{Z}_{\mathbf{A}}$ can be achieved following similar arguments.

Distributed Memory Implementation




Figure: Tensor partition.

- We assume p processing elements [2].

- We partition the matricization $\mathbf{X}_{\mathbf{A}}$ into p block rows as $\mathbf{X}_{\mathbf{A}} = \begin{bmatrix} (\mathbf{X}_{\mathbf{A}}^1)^T \\ \vdots \\ (\mathbf{X}_{\mathbf{A}}^p)^T \end{bmatrix}^T$, with

$\mathbf{X}_{\mathbf{A}}^n \in \mathbb{R}_+^{I \times JK}$. We partition similarly $\mathbf{X}_{\mathbf{B}}$ and $\mathbf{X}_{\mathbf{C}}$.

- The n -th block row of $\mathbf{X}_{\mathbf{A}}$, $\mathbf{X}_{\mathbf{B}}$, $\mathbf{X}_{\mathbf{C}}$ have been allocated to the n -th processing element, for $n = 1, \dots, p$.

- We partition \mathbf{A}_l into p block rows as $\mathbf{A}_l = \begin{bmatrix} (\mathbf{A}_l^1)^T \\ \vdots \\ (\mathbf{A}_l^p)^T \end{bmatrix}^T$, with $\mathbf{A}_l^n \in \mathbb{R}_+^{I \times R}$, for $n = 1, \dots, p$.

- The n -th processor knows the whole \mathbf{A}_l , but updates the n -th block row of \mathbf{A}_l , \mathbf{A}_l^n , for $n = 1, \dots, p$.

Factor Update Implementation

The update of \mathbf{A}_l is achieved via the updates of \mathbf{A}_l^n , for $n = 1, \dots, p$:

- The n -th processing element uses its local data $\mathbf{X}_{\mathbf{A}}^n$, as well as the whole matrices \mathbf{B}_l and \mathbf{C}_l , and computes the n -th block row of matrix \mathbf{A}_{l+1} , \mathbf{A}_{l+1}^n , via the Nesterov Matrix Completion algorithm.

- Each processing element broadcasts its output to all other processing elements; this operation can be implemented via an Allgather operation.

At this point, all processors know \mathbf{A}_{l+1} and are ready for the update of \mathbf{B}_l (and, then, of \mathbf{C}_l).

Numerical Experiments

- Results obtained from a Message Passing Interface (MPI) implementation of the AO NTC.
- The program is executed on a DELL PowerEdge R820 system with SandyBridge - Intel(R) Xeon(R) CPU E5 - 4650v2 (in total, 16 nodes with 40 cores each at 2.4 Gz) and 512 GB RAM per node.
- The matrix operations are implemented using routines of the C++ library Eigen [3].
- The performance metric we compute is the speedup attained using p processors.

Real Data

- The MovieLens 10M dataset [4], which contains time-stamped ratings of movies.
 - Binning the time into seven-day-wide bins, results in a tensor of size $71567 \times 65133 \times 171$.
 - The number of samples is 8M (99.99% sparsity).
 - We first perform a random permutation on our data to resolve load imbalance issues.
 - We measure the completion accuracy by measuring the mean squared error of 2M known ratings with our predictions. The mean squared error we achieved is 0.0033
- (For the n -th known rating, with indices (i_n, j_n, k_n) , we compute our prediction after rounding the quantity $\sum_{r=1}^R \mathbf{A}(i_n, r) \circledast \mathbf{B}(j_n, r) \circledast \mathbf{C}(k_n, r)$ to the closest integer.)

Speed Up

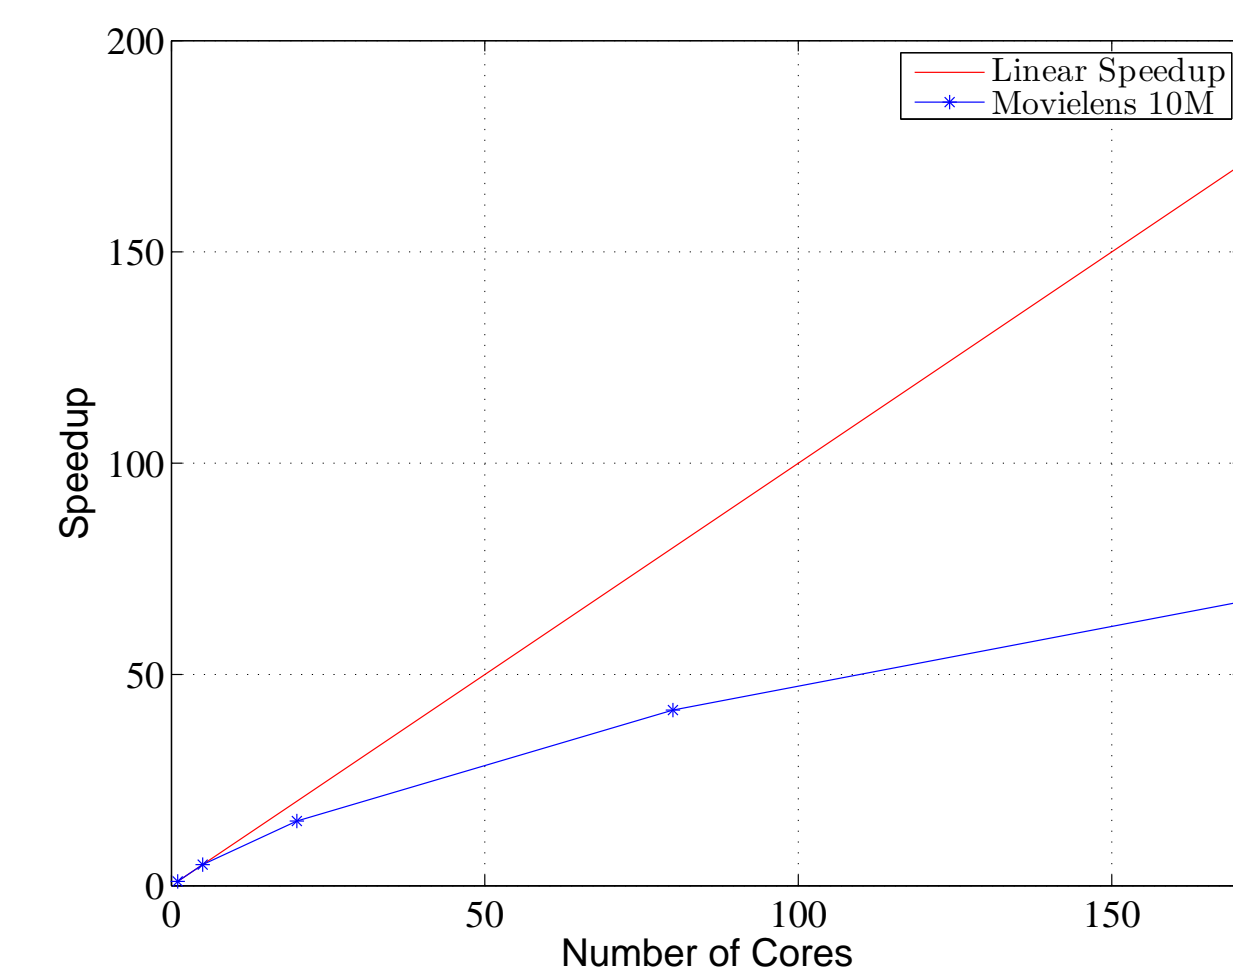


Figure: Speedup achieved for the MovieLens 10M dataset of size $71567 \times 65133 \times 171$ with p cores, for $p = 1, 5, 20, 171$.

Synthetic Data

- Synthetic data of the same size and sparsity level.
- True latent factors with i.i.d elements, uniformly distributed in $[0, 1]$.

Speed Up

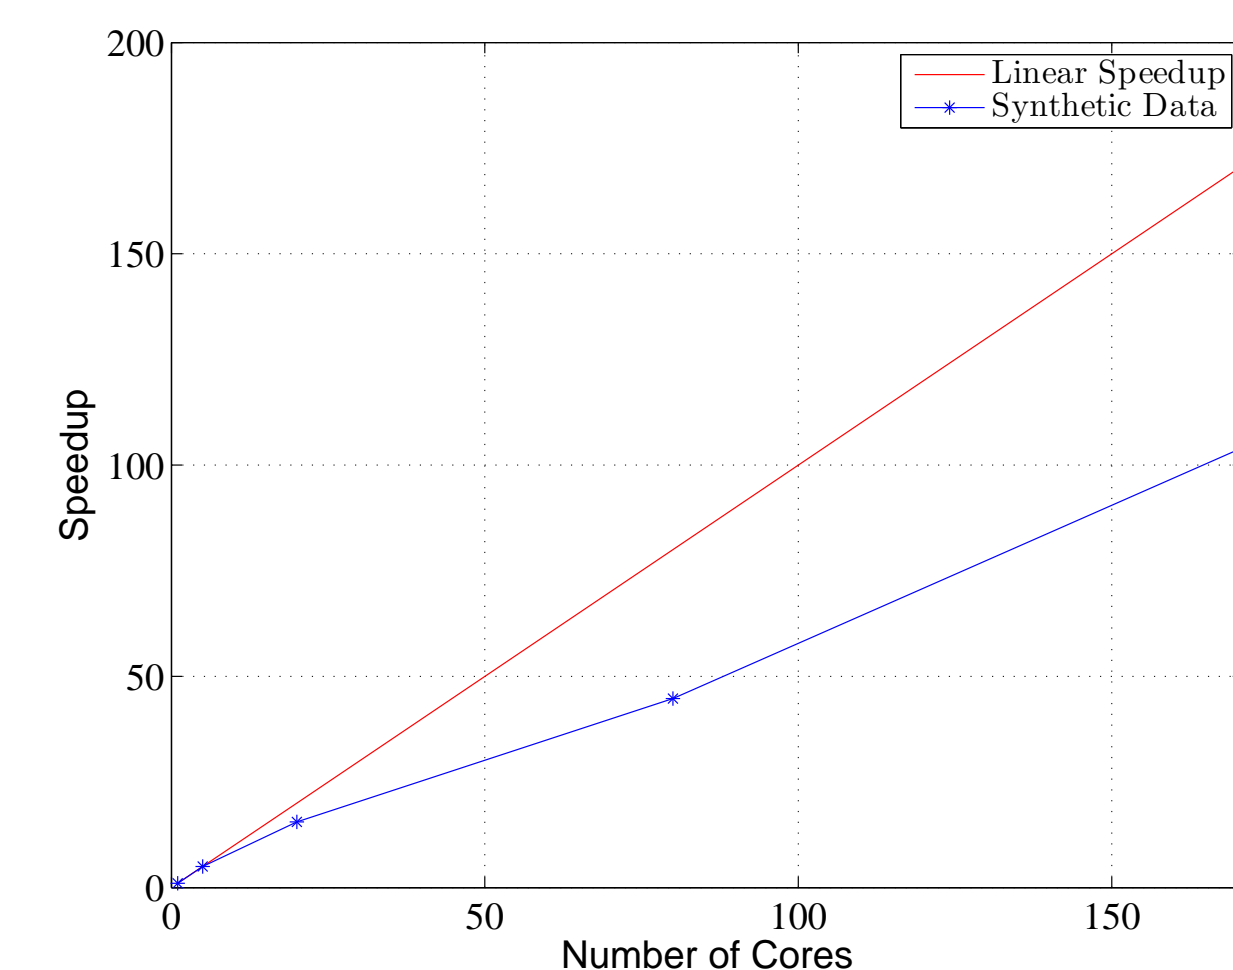


Figure: Speedup achieved for a $71567 \times 65133 \times 171$ tensor with p cores, for $p = 1, 5, 20, 171$.

References

- L. Karlsson, D. Kressner and A. Uschmajew, "Parallel algorithms for tensor completion in the CP format," *Parallel Computing*, 2015.
- K. Shin U. Kang, "Distributed Methods for High-dimensional and Large-scale Tensor Factorization," *IEEE International Conference on Data Mining*, pp. 989–994, 2014.
- G. Guennebaud and B. Jacob *et al.*, "Eigen v3," <http://eigen.tuxfamily.org>, 2010.
- F. Maxwell Harper and J. A. Konstan, "The MovieLens Datasets: History and Context," *ACM Transactions on Interactive Intelligent Systems (TiiS)*, vol. 5, no. 4, pp. 1–19, June, 2015.