Non-convex Approaches for Low-rank Tensor Completion under Tubal Sampling

Authors: Zheng Tan, Longxiu Huang, HanQin Cai, Yifei Lou

University of California, Los Angeles; Michigan State University; University of Central Florida; University of Texas at Dallas.

Abstract

Tensor completion is an important problem in modern data analysis. In this work, we investigate a specific sampling strategy, referred to as tubal sampling. We propose two novel non-convex tensor completion frameworks that are easy to implement, named tensor L1-L2 (TL12) and tensor completion via CUR (TCCUR). We test the efficiency of both methods on synthetic data and a color image inpainting problem. Empirical results reveal a trade-off between the accuracy and time efficiency of these two methods in a low sampling ratio. Each of them outperforms some classical completion methods in at least one aspect.

Introduction

This paper considers a tensor completion problem when each tubal is either sampled entirely or not sampled at all, which is referred to as tensor tubal sampling [5], as illustrated in Fig 1 (left).

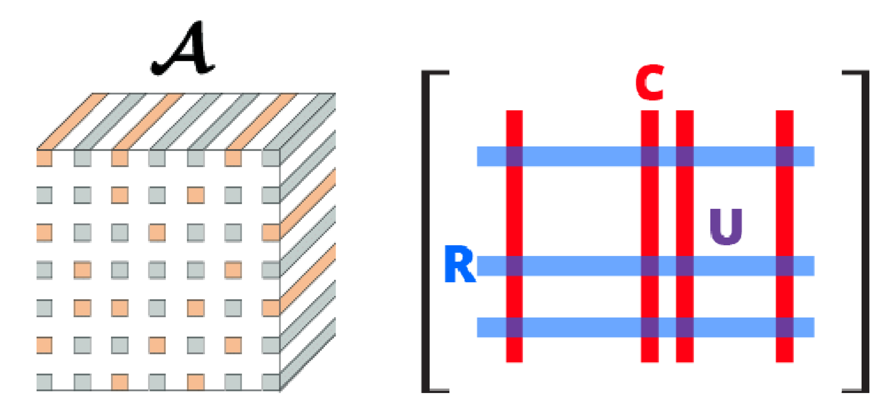

Fig. 1: Tensor Tubal Sampling (left) and matrix CUR (right).

Tensor, a multidimensional generalization of the matrix, is a useful data structure that is arisen in various fields. The matrix's SVD can be extended into tensors, called t-SVD [5]. Specifically for a general tensor $\mathcal{T} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$, the t-SVD is visualized in Fig. 2.

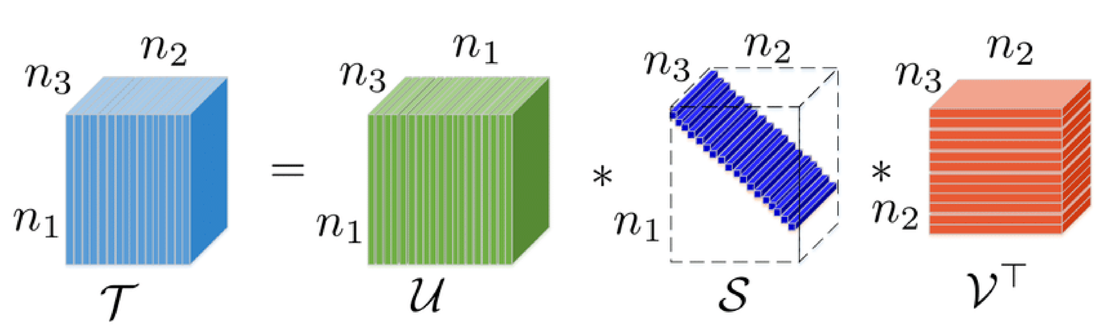

Fig. 2: Tensor SVD (t-SVD).

Tensor completion is an ill-posed problem, thus requiring additional information to be imposed as a regularization. We focus on a low-rank structure of the desired tensor, which requires defining the tensor's tubal rank as follows,

Definition 1: (Tensor multi rank and tubal rank) Let $\hat{\mathcal{A}} = \mathrm{fft}(\mathcal{A}, [], 3)$ to be the Fourier transform of $\mathcal{A}$ along the third dimension. The tensor multi rank of a 3-mode tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is a vector $\mathbf{r} \in \mathbb{R}^{n_3}$ with its i-th component equal to the rank of the i-th frontal slice of $\hat{\mathcal{A}}$. The tensor tubal rank of $\mathcal{A}$ is defined to be $r = \|\mathbf{r}\|_\infty$.

Objectives of this project:

(1) Propose two novel non-convex low-rank tensor completion under tubal sampling

(2) Draw empirical guidance in real applications.

Tensor Low-rank Regularization

We consider a general model for low-rank tensor completion

$$\min_{\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times n_3}} h(\mathcal{X}) \quad s.t. \quad \mathcal{Y} = \mathcal{P}_\Omega(\mathcal{X}),$$

where $h(\cdot)$ is a low-rank regularization and $\mathcal{P}_\Omega$ is a sampling operator

$$[\mathcal{P}_\Omega(\mathcal{X})]_{i,j,k} := \begin{cases} [\mathcal{X}]_{i,j,k}, & \text{if } (i,j,k) \in \Omega \\ 0, & \text{Otherwise.} \end{cases}$$

We propose the tensor L_1-L_2 (TL12) regularization,

$$||\mathcal{X}||_{\mathrm{TL12}} = \sum_{j=1}^{n_3} \left(\sum_{i=1}^{m} [\mathcal{S}]_{i,i,j} - \sqrt{\sum_{i=1}^{m} [\mathcal{S}]_{i,i,j}^2} \right).$$

By defining a vector $\mathbf{s}_j = ([\mathcal{S}]_{i,i,j})_{i=1,\cdots,m}$, the TL12 regularization is equivalent to the difference between the L_1 and L_2 norms of $\mathbf{s}_j$, followed by summing over $j = 1, \ldots, n_3$. The classic tensor nuclear norm (TNN) is defined by $||\mathcal{X}||_{\mathrm{TNN}} = \sum_{j=1}^{n_3} \mathcal{S}]_{i,i,j}$. Numerically, we adopt the alternating direction method of multipliers (ADMM) that iterates as follows,

$$\begin{aligned} \mathcal{X}^{(\ell+1)} &= \arg\min_{\mathcal{X}} \{||\mathcal{X} - (\mathcal{Z}^{(\ell)} - \mathcal{B}^{(\ell)})||_{\mathrm{F}}^2 \text{ s.t. } \mathcal{Y} = \mathcal{P}_\Omega(\mathcal{X})\} \\ \mathcal{Z}^{(\ell+1)} &= \arg\min_{\mathcal{Z}} \{\frac{1}{\rho} h(\mathcal{Z}) + \frac{1}{2} ||\mathcal{Z} - (\mathcal{X}^{(\ell+1)} + \mathcal{B}^{(\ell)})||_{\mathrm{F}}^2\} \\ \mathcal{B}^{(\ell+1)} &= \mathcal{B}^{(\ell)} + (\mathcal{X}^{(\ell)} - \mathcal{Z}^{(\ell+1)}), \end{aligned}$$

where $\mathcal{Z}$ is an auxiliary variable, $\mathcal{B}$ is a Lagrangian multiplier to enforce $\mathcal{X} = \mathcal{Z}$, $\rho > 0$ is a weighting parameter, and ℓ counts the iterations. The algorithm alternates between $\mathcal{X}$ satisfying the data matching constraint and promoting $\mathcal{Z}$ to be low-rank. The closed-form solution for $\mathcal{X}^{(\ell+1)}$ is that it takes the values of $\mathcal{Y}$ on Ω and of $\mathcal{Z}^{(\ell)} - \mathcal{B}^{(\ell)}$ on the complement set of Ω. The $\mathcal{Z}$-subproblem has a closed-form solution based on the proximal operator of L_1-L_2 [3].

Tensor Completion via CUR

By the design of tubal sampling, we have $\hat{\mathcal{Y}} = \mathcal{P}_\Omega(\hat{\mathcal{X}})$. As a result, we can find an estimate of $\mathcal{X}$ by completing $\hat{\mathcal{Y}}$, followed by the inverse Fourier transform along the third dimension. Completing the tensor $\hat{\mathcal{Y}}$ reduces to a series of matrix completion problems, independently for each frontal slice of $\hat{\mathcal{Y}}$. We adopt a recently developed matrix completion method termed iterative CUR completion (ICURC) [1]. THe matrix CUR is illustrated in Fig 1 (right). The definition of tensor-CUR (t-CUR) is given in Definition 2.

Definition 2: (t-CUR) The t-CUR decomposition of $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is given by $\mathcal{C} * \mathcal{U}^\dagger * \mathcal{R}$, where $\mathcal{C} = [\mathcal{A}]_{:,J,:}$, $\mathcal{R} = [\mathcal{A}]_{I,:,:}$, $\mathcal{U} = [\mathcal{A}]_{I,J,:}$ with $I \subseteq [n_1]$ and $J \subseteq [n_2]$.

Starting with $X^{(0)} = 0$ and the CUR decomposition of $X^{(\ell)} = C^{(\ell)}(U^{(\ell)})^\dagger R^{(\ell)}$ at the ℓth iteration, we advance to the next step by

$$\begin{aligned} [C^{(\ell+1)}]_{I^c,:} &= [X^{(\ell)}]_{I^c,J} + [Y - \mathcal{P}_\Phi(X^{(\ell)})]_{I^c,J}, \\ [R^{(\ell+1)}]_{:,J^c} &= [X^{(\ell)}]_{I,J^c} + [Y - \mathcal{P}_\Phi(X^{(\ell)})]_{I,J^c}, \end{aligned}$$

where $I^c = [n_1] \setminus I$ and $J^c = [n_2] \setminus J$. The update of $U^{(\ell+1)}$ requires the best rank r approximation, which can be achieved by truncating the largest r singular values in the matrix SVD, denoted by $\mathcal{H}_r$. In short, we have the formula, $U^{(\ell+1)} = \mathcal{H}_r\left([X^{(\ell)}]_{I,J} + [Y - \mathcal{P}_\Phi(X^{(\ell)})]_{I,J}\right)$. With proper stopping conditions, we obtain the row and column submatrices $[C^{(\ell+1)}]_{\mathcal{I},:}$, $[R^{(\ell+1)}]_{:,\mathcal{J}}$, and $\mathcal{H}_r([X^{(\ell+1)}]_{\mathcal{I},\mathcal{J}})$.

Synthetic Experiments

We compare the performance of the two regularizations (TNN and TL12) and one decomposition method (TCCUR). We generate the sampling set $\Phi \subseteq [n_1] \times [n_2]$ uniformly at random under a preset sampling ratio (without replacements) to define the index set Ω. We evaluate the performance by the relative error (RE) and peak signal-to-noise ratio (PSNR), i.e.,

$$\mathrm{RE} = \frac{\|\mathcal{X} - \tilde{\mathcal{X}}\|_{\mathrm{F}}}{\|\mathcal{X}\|_{\mathrm{F}}} \text{ and } \mathrm{PSNR} = 10 \log_{10} \left(\frac{n_1 n_2 n_3 \mathcal{X}_{\max}^2}{\|\tilde{\mathcal{X}} - \mathcal{X}\|_{\mathrm{F}}^2} \right),$$

where $\tilde{\mathcal{X}} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is the recovered tensor and $\mathcal{X}$ is the ground truth with its maximum absolute value, denoted by $\mathcal{X}_{\max}$.

We generate a tensor $\mathcal{X} \in \mathbb{R}^{256 \times 256 \times 50}$ with tubal rank $r \in \{3, 5\}$. For each preset rank, Fig. 3 shows the mean of REs over 30 random realizations with respect to sampling ratio (SR), showing that

- TL12 and TCCUR achieve comparable and even better performance than TNN;
- TL12 method has the fastest decay of RE for smaller SRs (e.g., 10% − 20%).

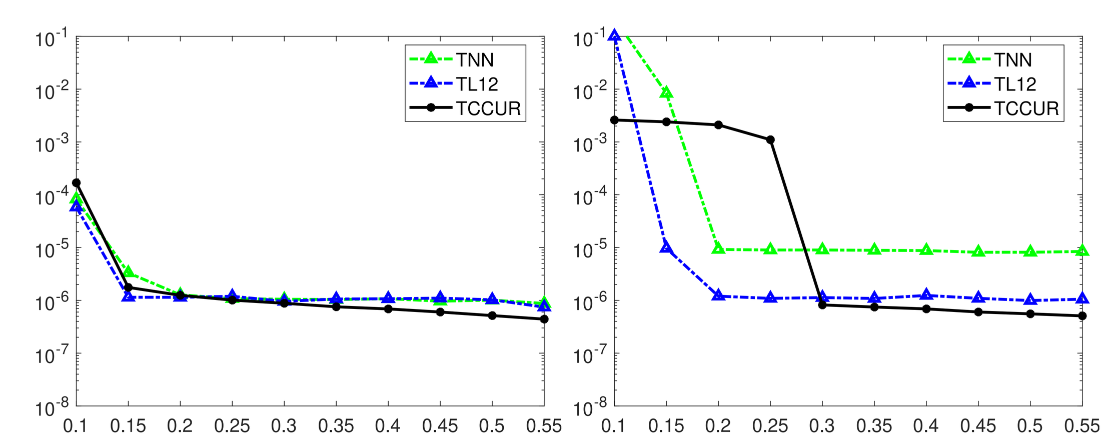

Fig. 3: REs of completing an underlying tensor of tubal rank 3 (left) and 5 (right) versus SRs.

We also examine the **scalability** of the algorithms by reporting the runtime with respect to the tensor's dimensions. In particular, we generate a tensor $\mathcal{X} \in \mathbb{R}^{2^n \times 2^n \times 32}$ with tubal rank 2 and 3 for $n = 6, 7, 8, 9, 10$. We randomly select 30% tubals and adopt the same stopping condition for all the algorithms, that is, the relative error on the observed portion is less than 10^{-6}. The computational time is reported in Fig. 4, illustrating

- significant advantages in the efficiency of TCCUR over TNN and TL12;
- TL12 is comparable in speed compared to TNN, yet gives better completion accuracy.

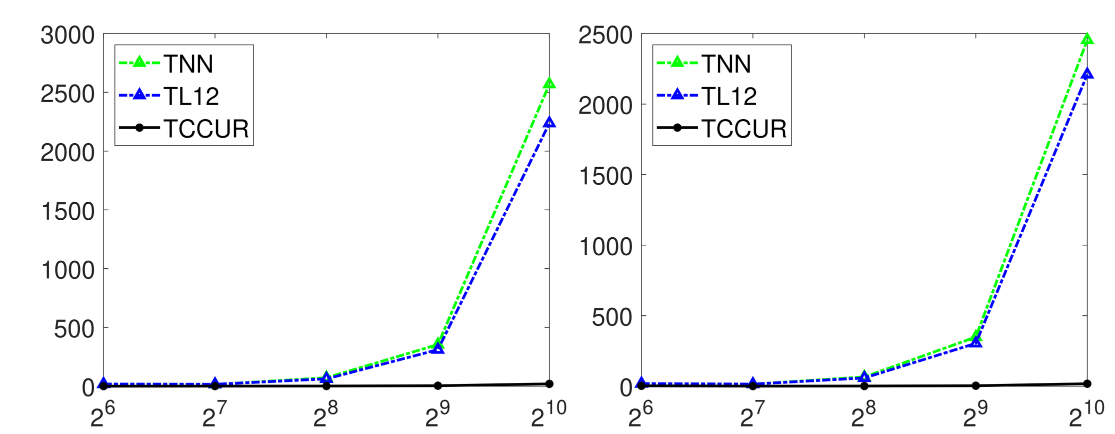

Fig. 4: Runtime versus the frontal slice dimension 2^n of $2^n \times 2^n \times 32$ tensors with tubal rank 2 (left) and 3 (right).

Image Inpainting Experiments

We investigate a real application of image inpainting on three color images used in [2]. We compare the proposed methods to two state-of-the-art methods named TRLRF [4] and PSTNN [2] with 50% randomly sampled tubals in Table 1.

- TCCUR is significantly faster than others;
- TL12 yields the best results in all test cases both visually and in terms of SNR, though slower than other methods

	Door		Hat		Starfish	
	PSNR	Time	PSNR	Time	PSNR	Time
TRLRF	30.01	5.02s	26.55	4.54s	23.93	13.85s
PSTNN	28.13	13.58s	19.38	2.34s	16.08	4.68s
TL12	31.13	63.60s	27.12	8.59s	25.94	22.08s
TCCUR	28.27	4.96s	26.37	1.22s	24.14	4.27s

Table 1: Comparison of image inpainting from 50% tubal sampling ratio.

Fig. 5 shows the reconstruction results; PSTNN clearly fails in filling reasonable values, while TRLRF and TCCUR produce more severe artifacts near the rim of the hat, compared to TL12.

TRLRF PSTNN TL12 TCCUR

Fig. 5: Visual comparison of color image inpainting results.

Conclusions

- The regularization-based method (TL12) achieves high accuracy in tensor completion but at a cost of high computational complexity.
- The decomposition method (TCCUR) is efficient, but its usage is limited to tubal sampling.

References

[1] HanQin Cai et al. *Matrix Completion with Cross-Concentrated Sampling: Bridging Uniform Sampling and CUR Sampling*. 2022.

[2] Tai-Xiang Jiang et al. *Multi-dimensional imaging data recovery via minimizing the partial sum of tubal nuclear norm*. 2020.

[3] Yifei Lou and Ming Yan. *Fast L1−L2 Minimization via a Proximal Operator*. 2018.

[4] Longhao Yuan et al. *Tensor ring decomposition with rank minimization on latent space: An efficient approach for tensor completion*. 2019.

[5] Zemin Zhang and Shuchin Aeron. *Exact tensor completion using t-SVD*. 2016.